



(12) **United States Patent**
Sivasubramanian et al.

(10) **Patent No.:** **US 8,060,792 B2**
(45) **Date of Patent:** **Nov. 15, 2011**

(54) **MONITORING AND AUTOMATED RECOVERY OF DATA INSTANCES**

(75) Inventors: **Swaminathan Sivasubramanian**, Seattle, WA (US); **Grant Alexander MacDonald McAlister**, Seattle, WA (US); **Barry B. Hunter, Jr.**, Sammamish, WA (US); **Parikshit S. Pol**, Bellevue, WA (US)

(73) Assignee: **Amazon Technologies, Inc.**, Reno, NV (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 197 days.

(21) Appl. No.: **12/415,987**

(22) Filed: **Mar. 31, 2009**

(65) **Prior Publication Data**

US 2010/0251002 A1 Sep. 30, 2010

(51) **Int. Cl.**
G06F 11/00 (2006.01)

(52) **U.S. Cl.** **714/47.1**

(58) **Field of Classification Search** None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,555,404	A *	9/1996	Torbjørnsen et al.	1/1
6,205,465	B1	3/2001	Schoening et al.	
6,671,821	B1	12/2003	Castro et al.	
6,961,768	B2 *	11/2005	Davis et al.	709/224
6,981,135	B1	12/2005	Trask	
7,062,559	B2	6/2006	Yoshimura et al.	
7,315,826	B1 *	1/2008	Guheen et al.	705/7
7,418,484	B2	8/2008	Presley	
7,502,329	B2 *	3/2009	Li et al.	370/252
7,536,686	B2	5/2009	Tan et al.	

7,624,133	B1	11/2009	Ojalvo	
2003/0005091	A1 *	1/2003	Ullmann et al.	709/220
2003/0212775	A1	11/2003	Steele et al.	
2003/0212898	A1	11/2003	Steele et al.	
2004/0073676	A1	4/2004	Honma et al.	
2004/0078637	A1	4/2004	Fellin et al.	
2004/0148443	A1	7/2004	Achiwa	
2004/0163008	A1 *	8/2004	Kim	714/4
2004/0174823	A1	9/2004	Steele et al.	
2005/0027845	A1	2/2005	Secor et al.	
2005/0193245	A1	9/2005	Hayden et al.	
2005/0210128	A1 *	9/2005	Cannon et al.	709/224
2005/0216788	A1	9/2005	Mani-Meitav et al.	
2005/0243611	A1	11/2005	Lubbers et al.	
2006/0106675	A1	5/2006	Cohen et al.	
2006/0106774	A1	5/2006	Cohen et al.	
2007/0022122	A1	1/2007	Bahar et al.	

(Continued)

OTHER PUBLICATIONS

U.S. Appl. No. 12/418,478, filed Apr. 3, 2009, Sheth.

(Continued)

Primary Examiner — Michael Maskulinski

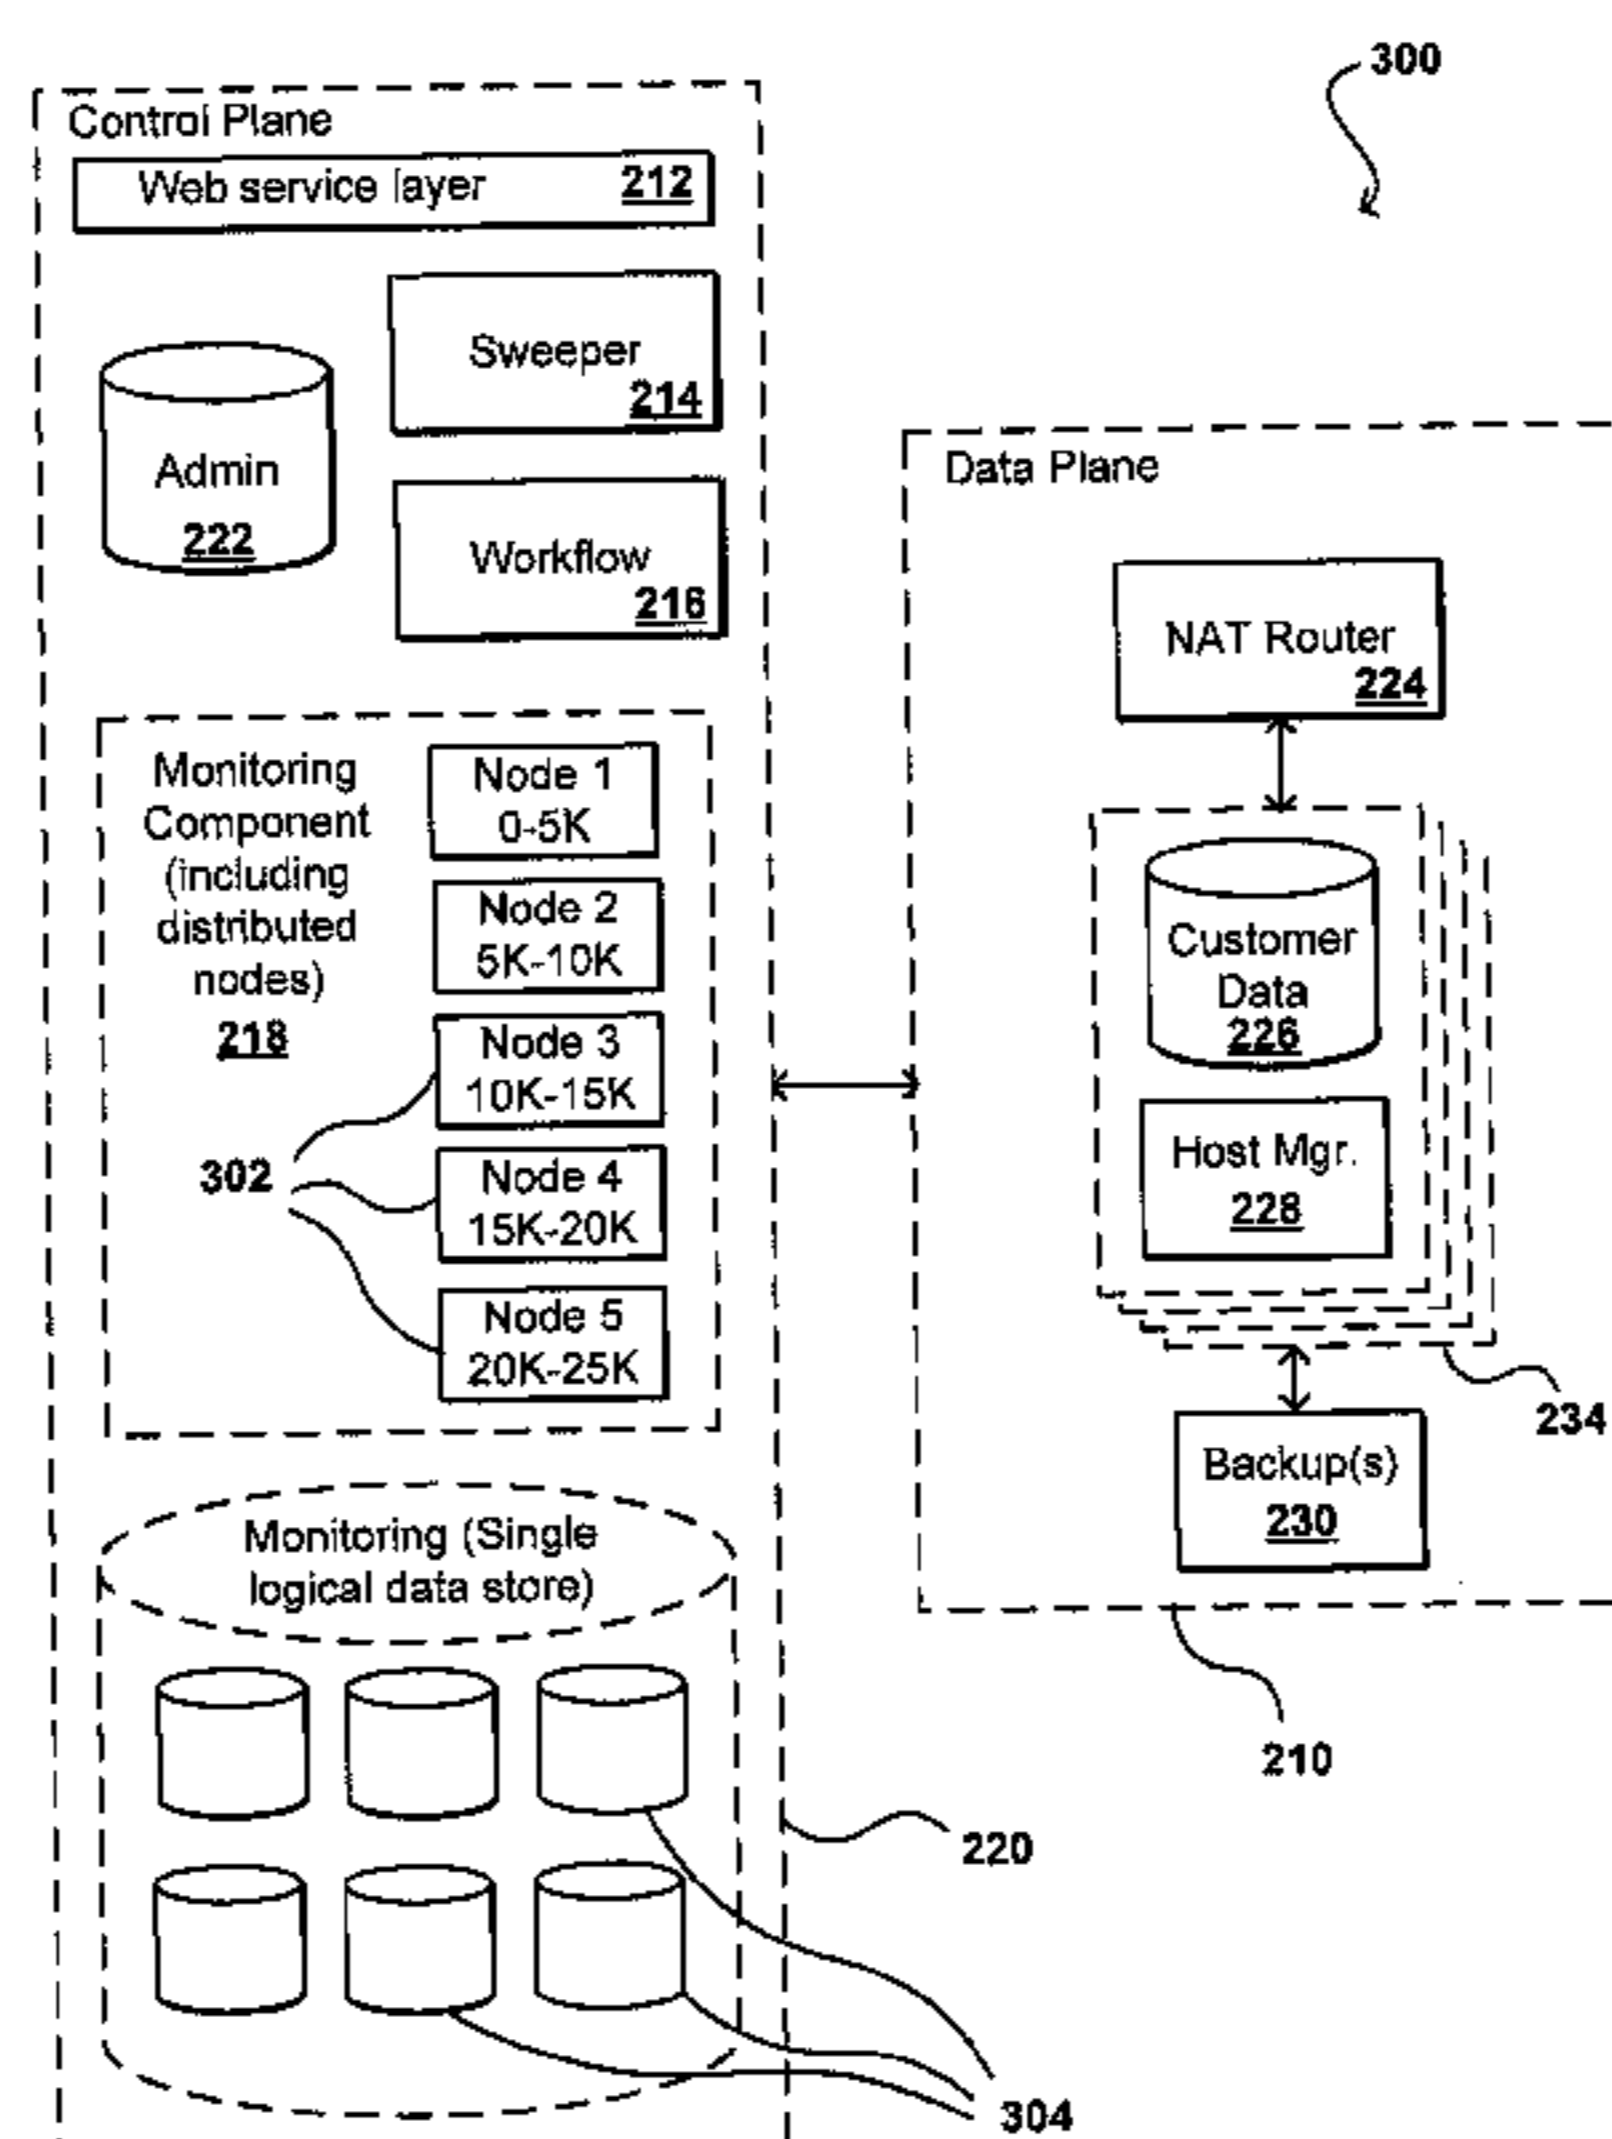
Assistant Examiner — Neil Miles

(74) *Attorney, Agent, or Firm* — Kilpatrick Townsend & Stockton LLP

(57) **ABSTRACT**

The monitoring and recovery of data instances, data stores, and other such components in a data environment can be performed automatically using a separate control environment. A monitoring component of the control plane can include a set of event processors for monitoring a workload of the data environment, where an event processor detecting a problem in the data plane can cause a recovery workflow to be generated in order to recover from the detected problem. The event processors can communicate with each other such that if one of the event processors becomes unavailable, the other event processors in a set are able to automatically redistribute responsibility for the workload.

19 Claims, 6 Drawing Sheets



U.S. PATENT DOCUMENTS

2007/0061266 A1 3/2007 Moore et al.
 2007/0061487 A1 3/2007 Moore et al.
 2007/0156842 A1 7/2007 Vermeulen et al.
 2007/0174691 A1 7/2007 D'Souza et al.
 2007/0244996 A1 10/2007 Ahmed et al.
 2007/0260696 A1 11/2007 Bohannon et al.
 2007/0288526 A1 12/2007 Mankad et al.
 2007/0299828 A1 12/2007 Lewis et al.
 2008/0065650 A1 3/2008 Kim et al.
 2008/0134176 A1 6/2008 Fitzgerald et al.
 2008/0134178 A1 6/2008 Fitzgerald et al.
 2008/0183991 A1 7/2008 Cosmadopoulos et al.
 2008/0189413 A1* 8/2008 Srivastava et al. 709/224
 2008/0263388 A1 10/2008 Allen et al.
 2009/0106411 A1 4/2009 Lisiecki et al.
 2009/0164853 A1 6/2009 Gokhale et al.
 2009/0198940 A1 8/2009 Ash et al.
 2009/0216881 A1* 8/2009 Lovy et al. 709/224
 2009/0271656 A1 10/2009 Yokota et al.
 2009/0328065 A1 12/2009 Wookey
 2010/0036851 A1 2/2010 Paterson-Jones et al.
 2010/0169707 A1* 7/2010 Mathew et al. 714/7
 2010/0191713 A1 7/2010 Lomet et al.
 2010/0250499 A1 9/2010 McAlister et al.
 2010/0250748 A1 9/2010 Sivasubramanian et al.

2010/0251242 A1 9/2010 Sivasubramanian et al.
 2010/0251339 A1 9/2010 McAlister
 2011/0004457 A1 1/2011 Haviv et al.
 2011/0083138 A1 4/2011 Sivasubramanian et al.
 2011/0099146 A1 4/2011 McAlister et al.
 2011/0099147 A1 4/2011 McAlister et al.
 2011/0099420 A1 4/2011 McAlister et al.

OTHER PUBLICATIONS

International Search Report mailed on May 25, 2010 for International Application No. PCT/US10/029476 filed on Mar. 31, 2010, 2 pages.
 International Search Report mailed on Dec. 2, 2010 for International Application No. PCT/US10/051757 filed on Oct. 7, 2010, 2 pages.
 International Search Report mailed on Dec. 21, 2010 for International Application No. PCT/US10/54133 filed on Oct. 26, 2010, 2 pages.
 International Search Report mailed on Dec. 21, 2010 for International Application No. PCT/US10/54139 filed on Oct. 26, 2010, 2 pages.
 International Search Report mailed on Dec. 21, 2010 for International Application No. PCT/US10/54141 filed on Oct. 26, 2010, 2 pages.

* cited by examiner




FIG. 1

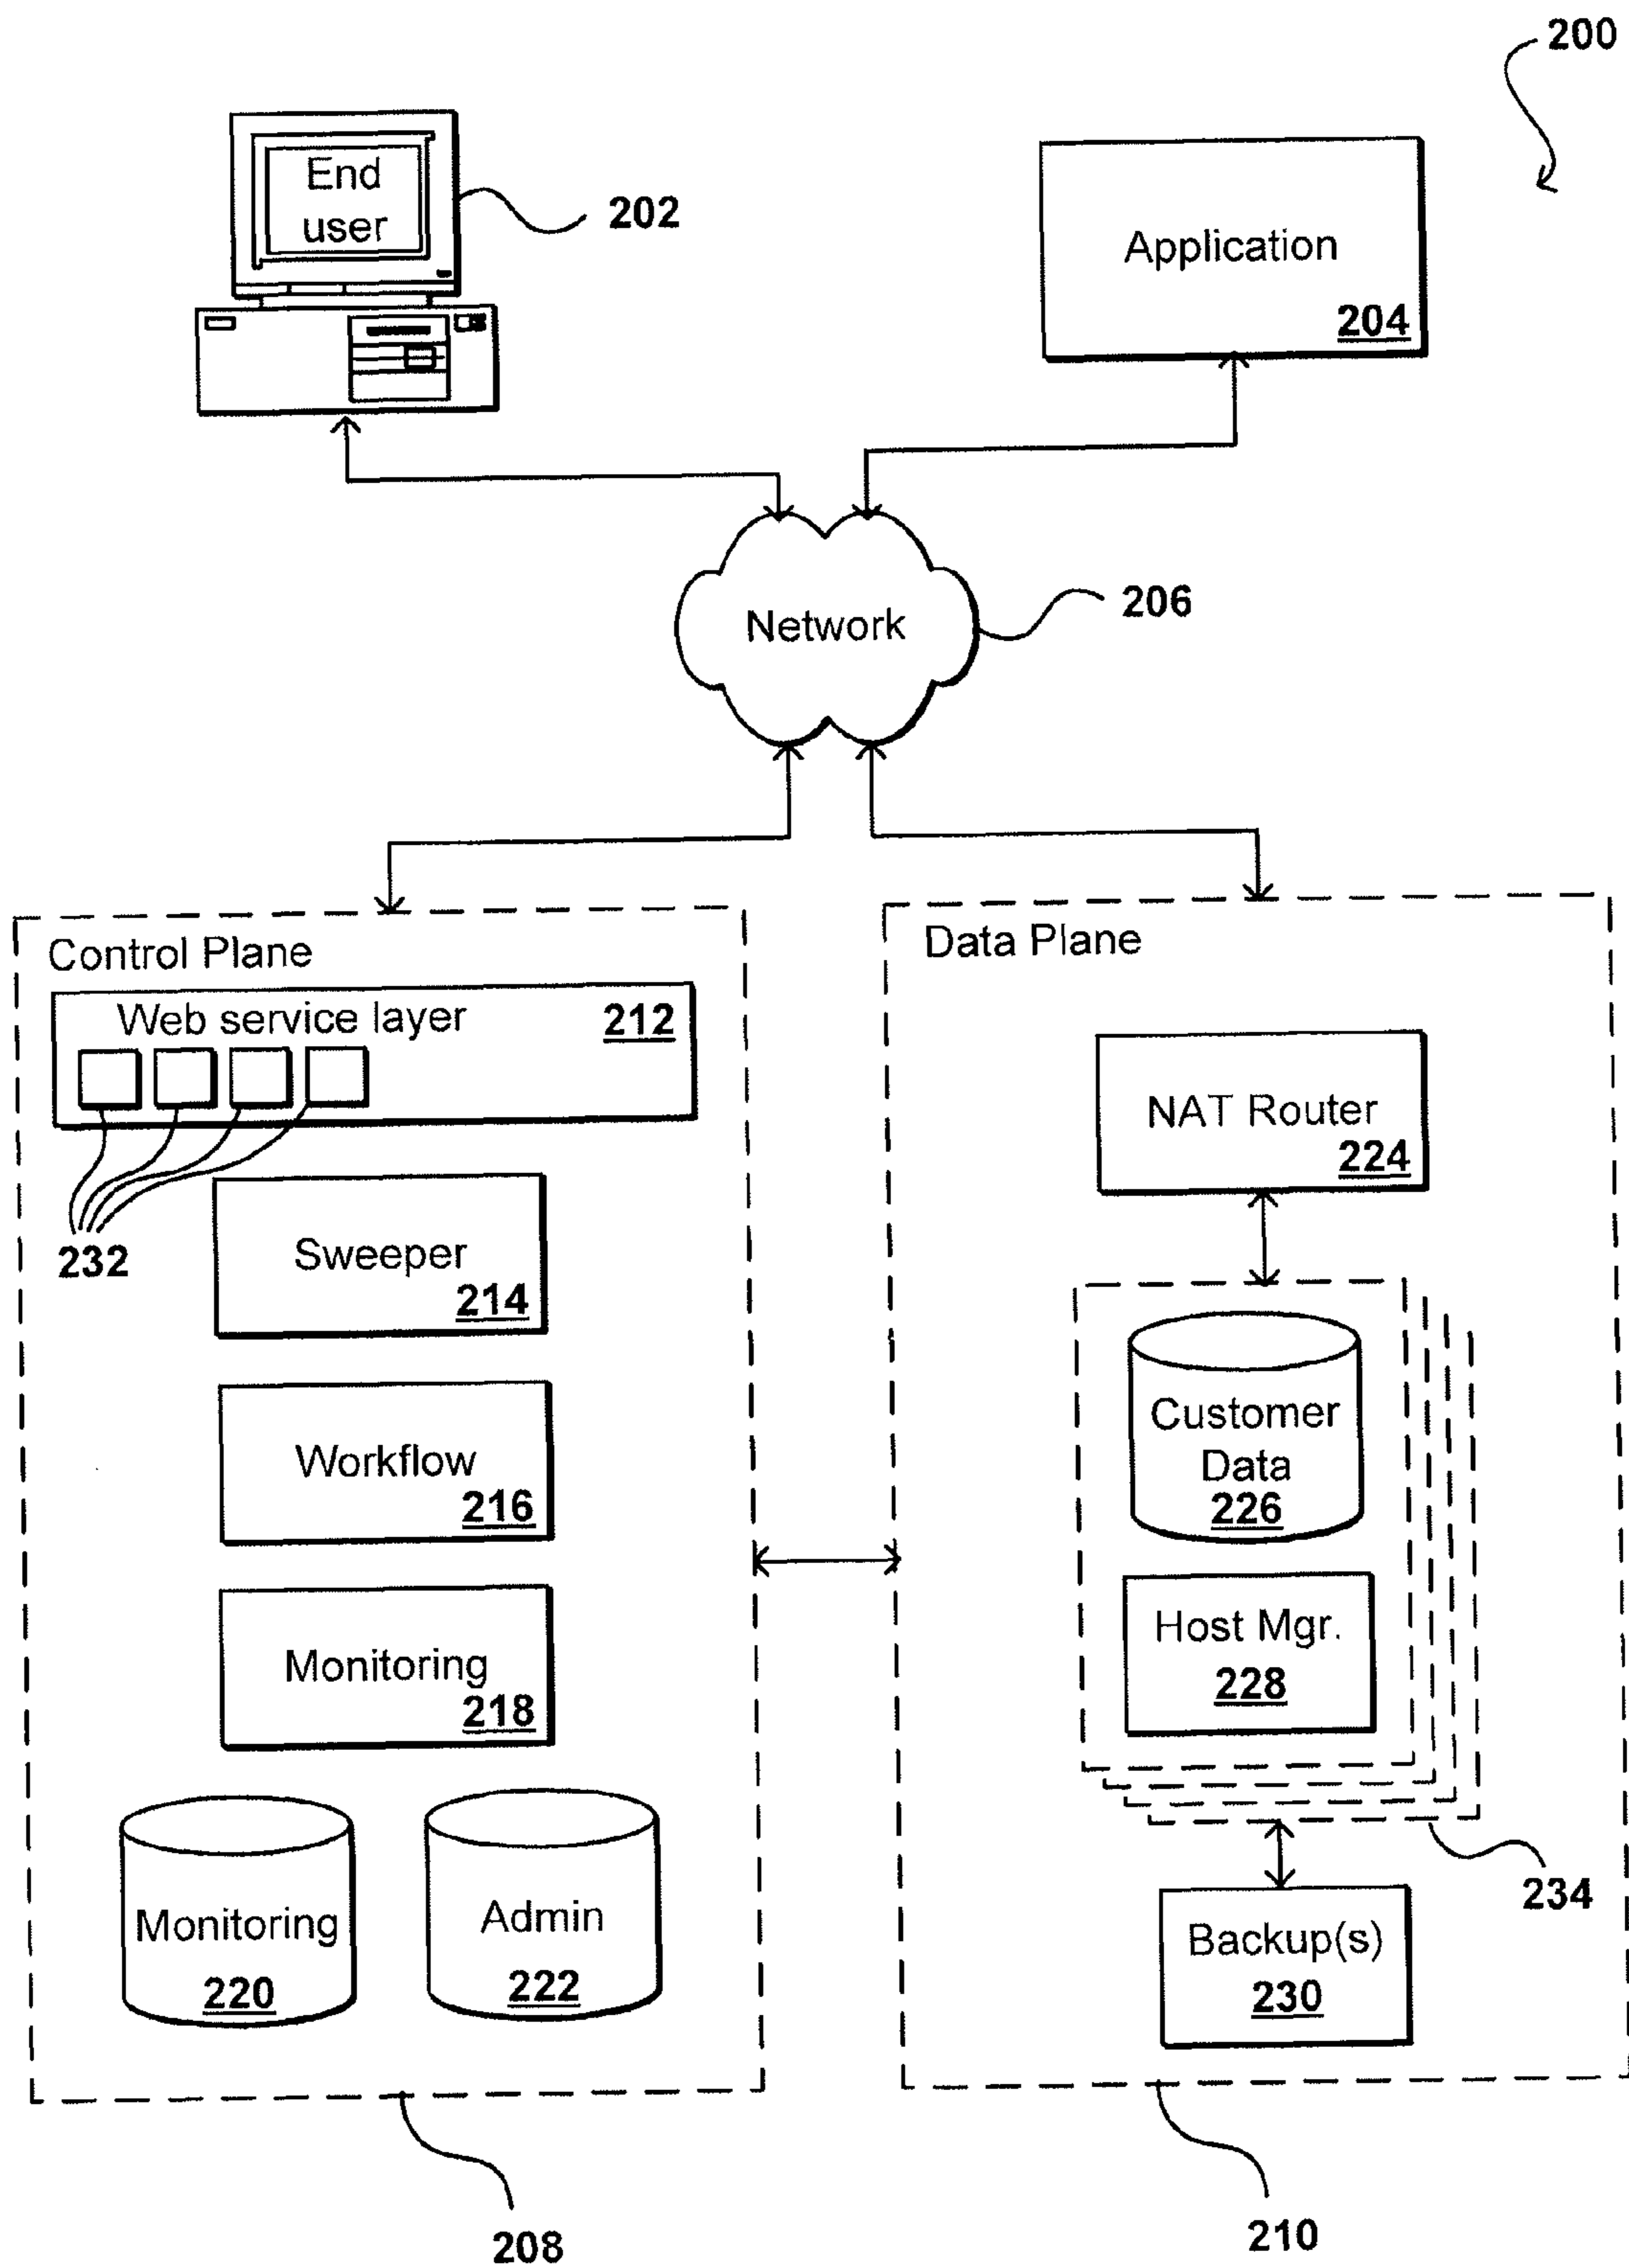


FIG. 2




FIG. 3

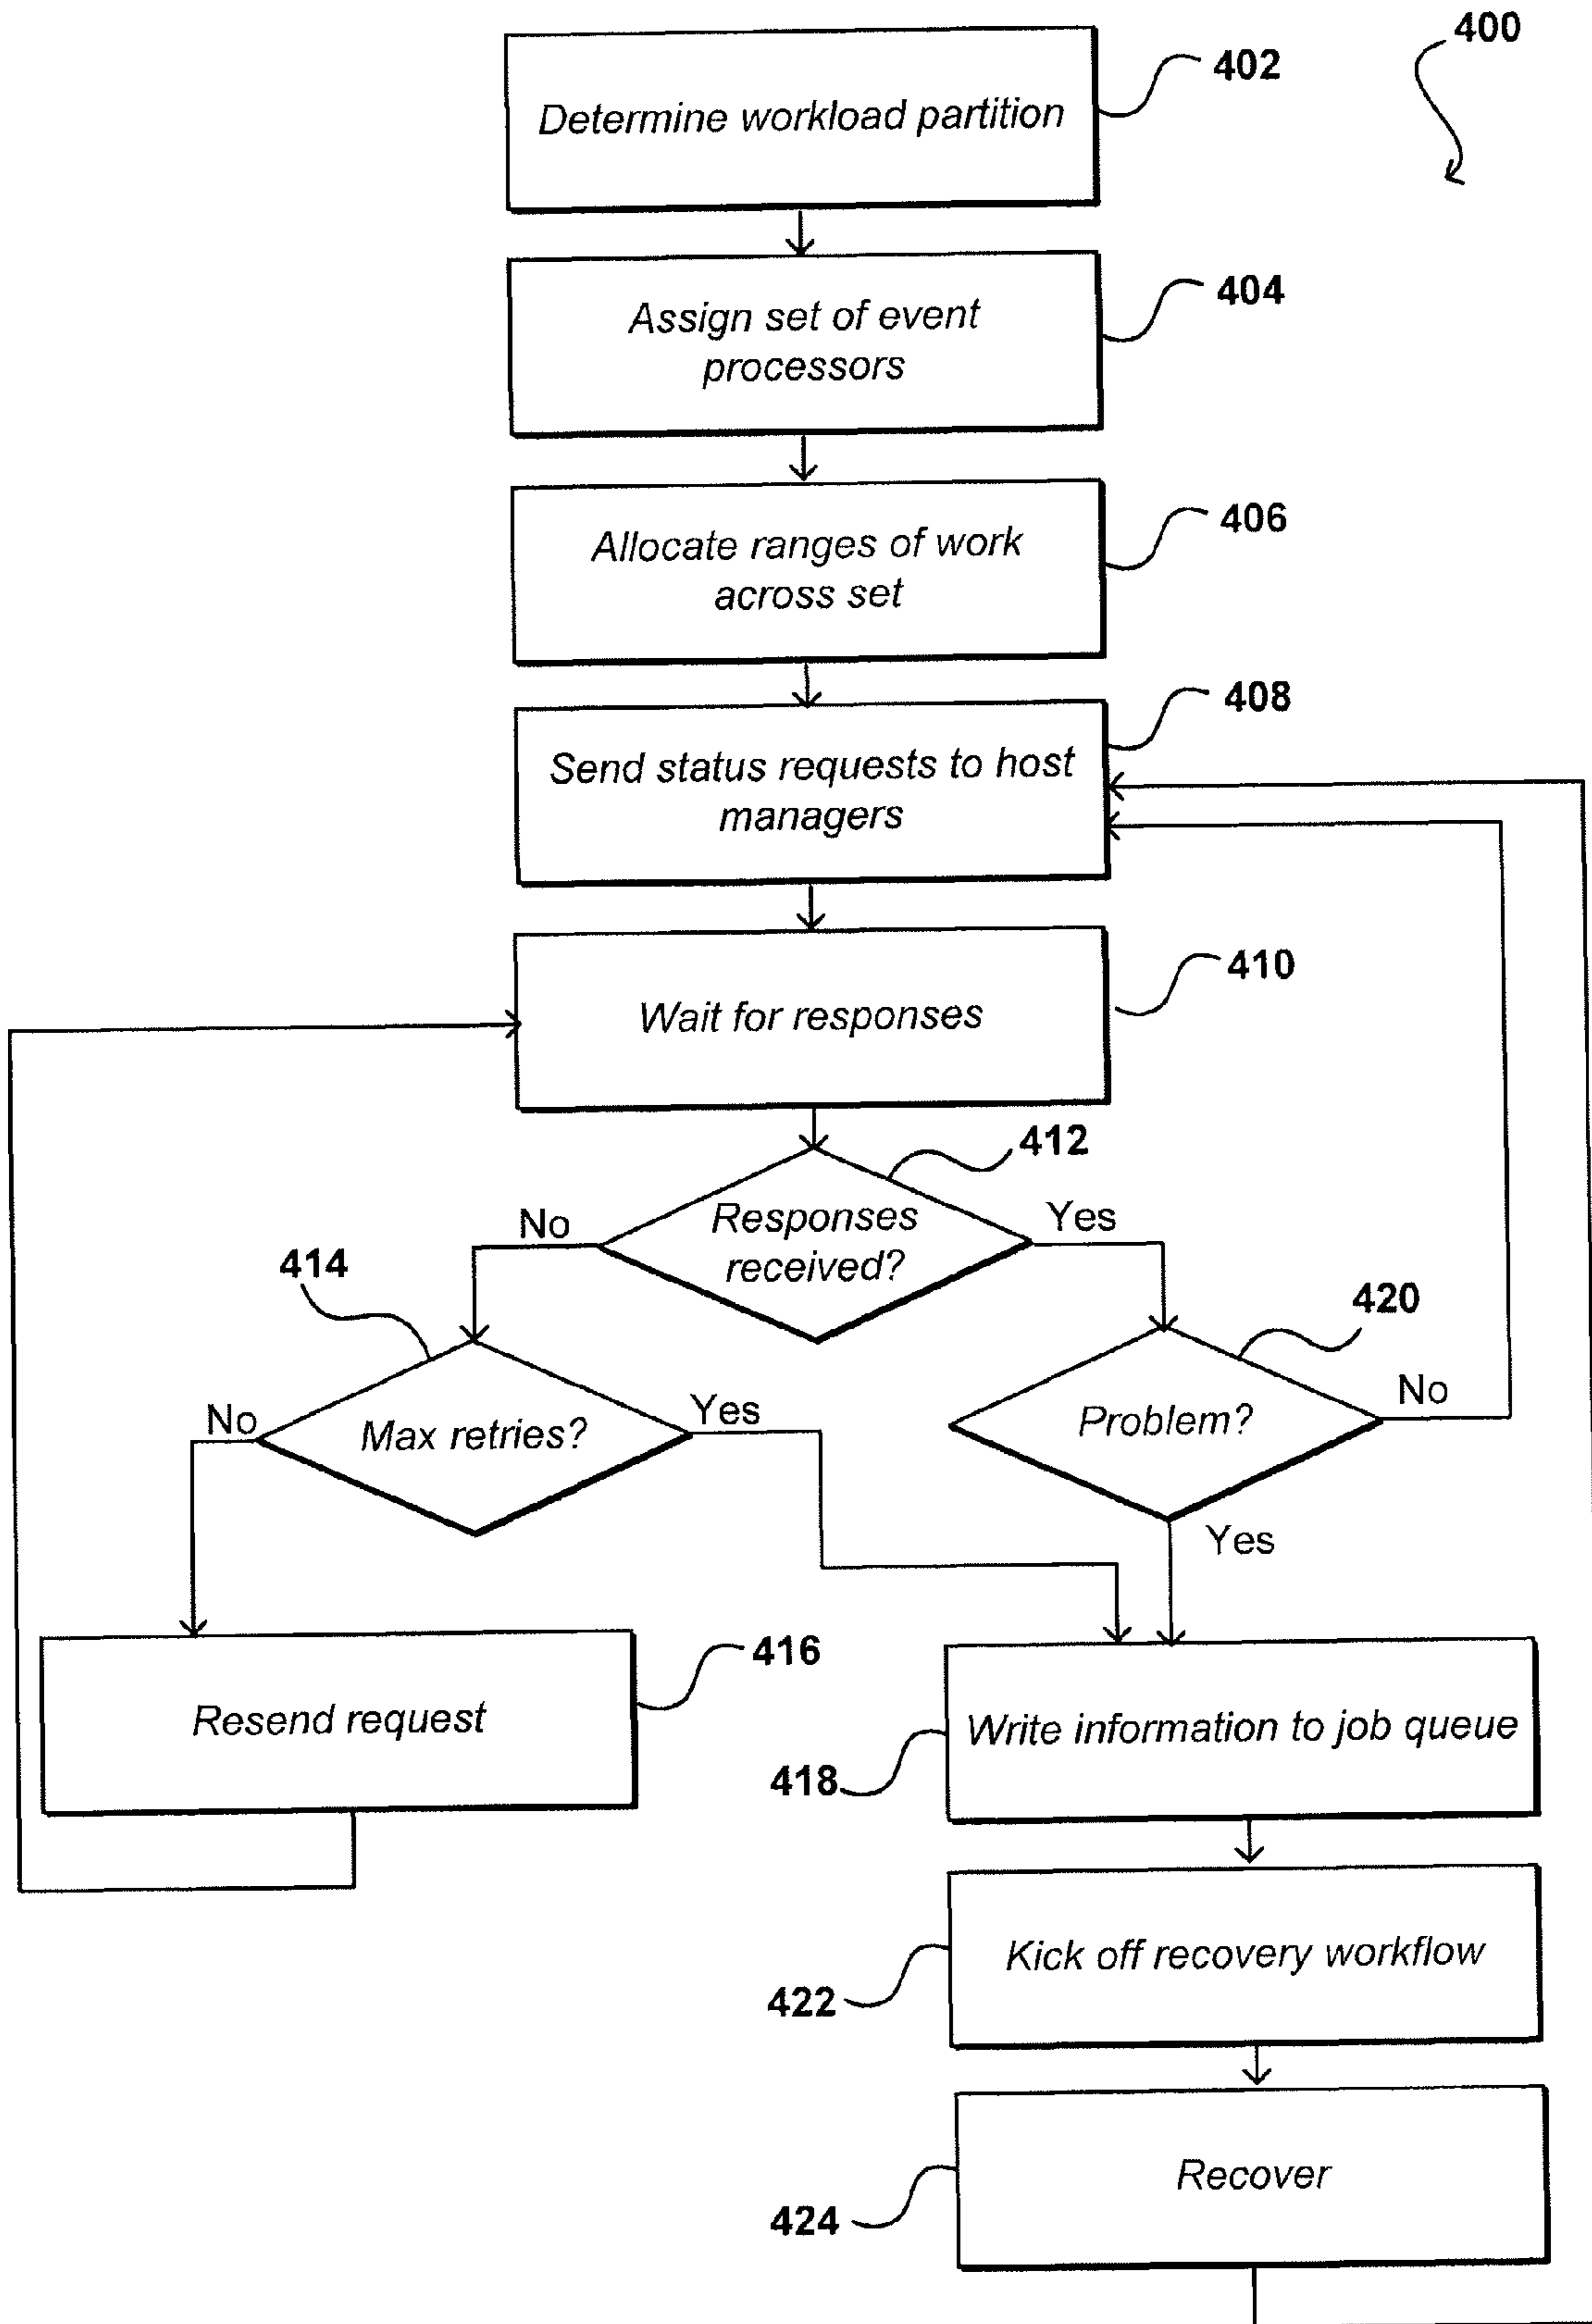


FIG. 4

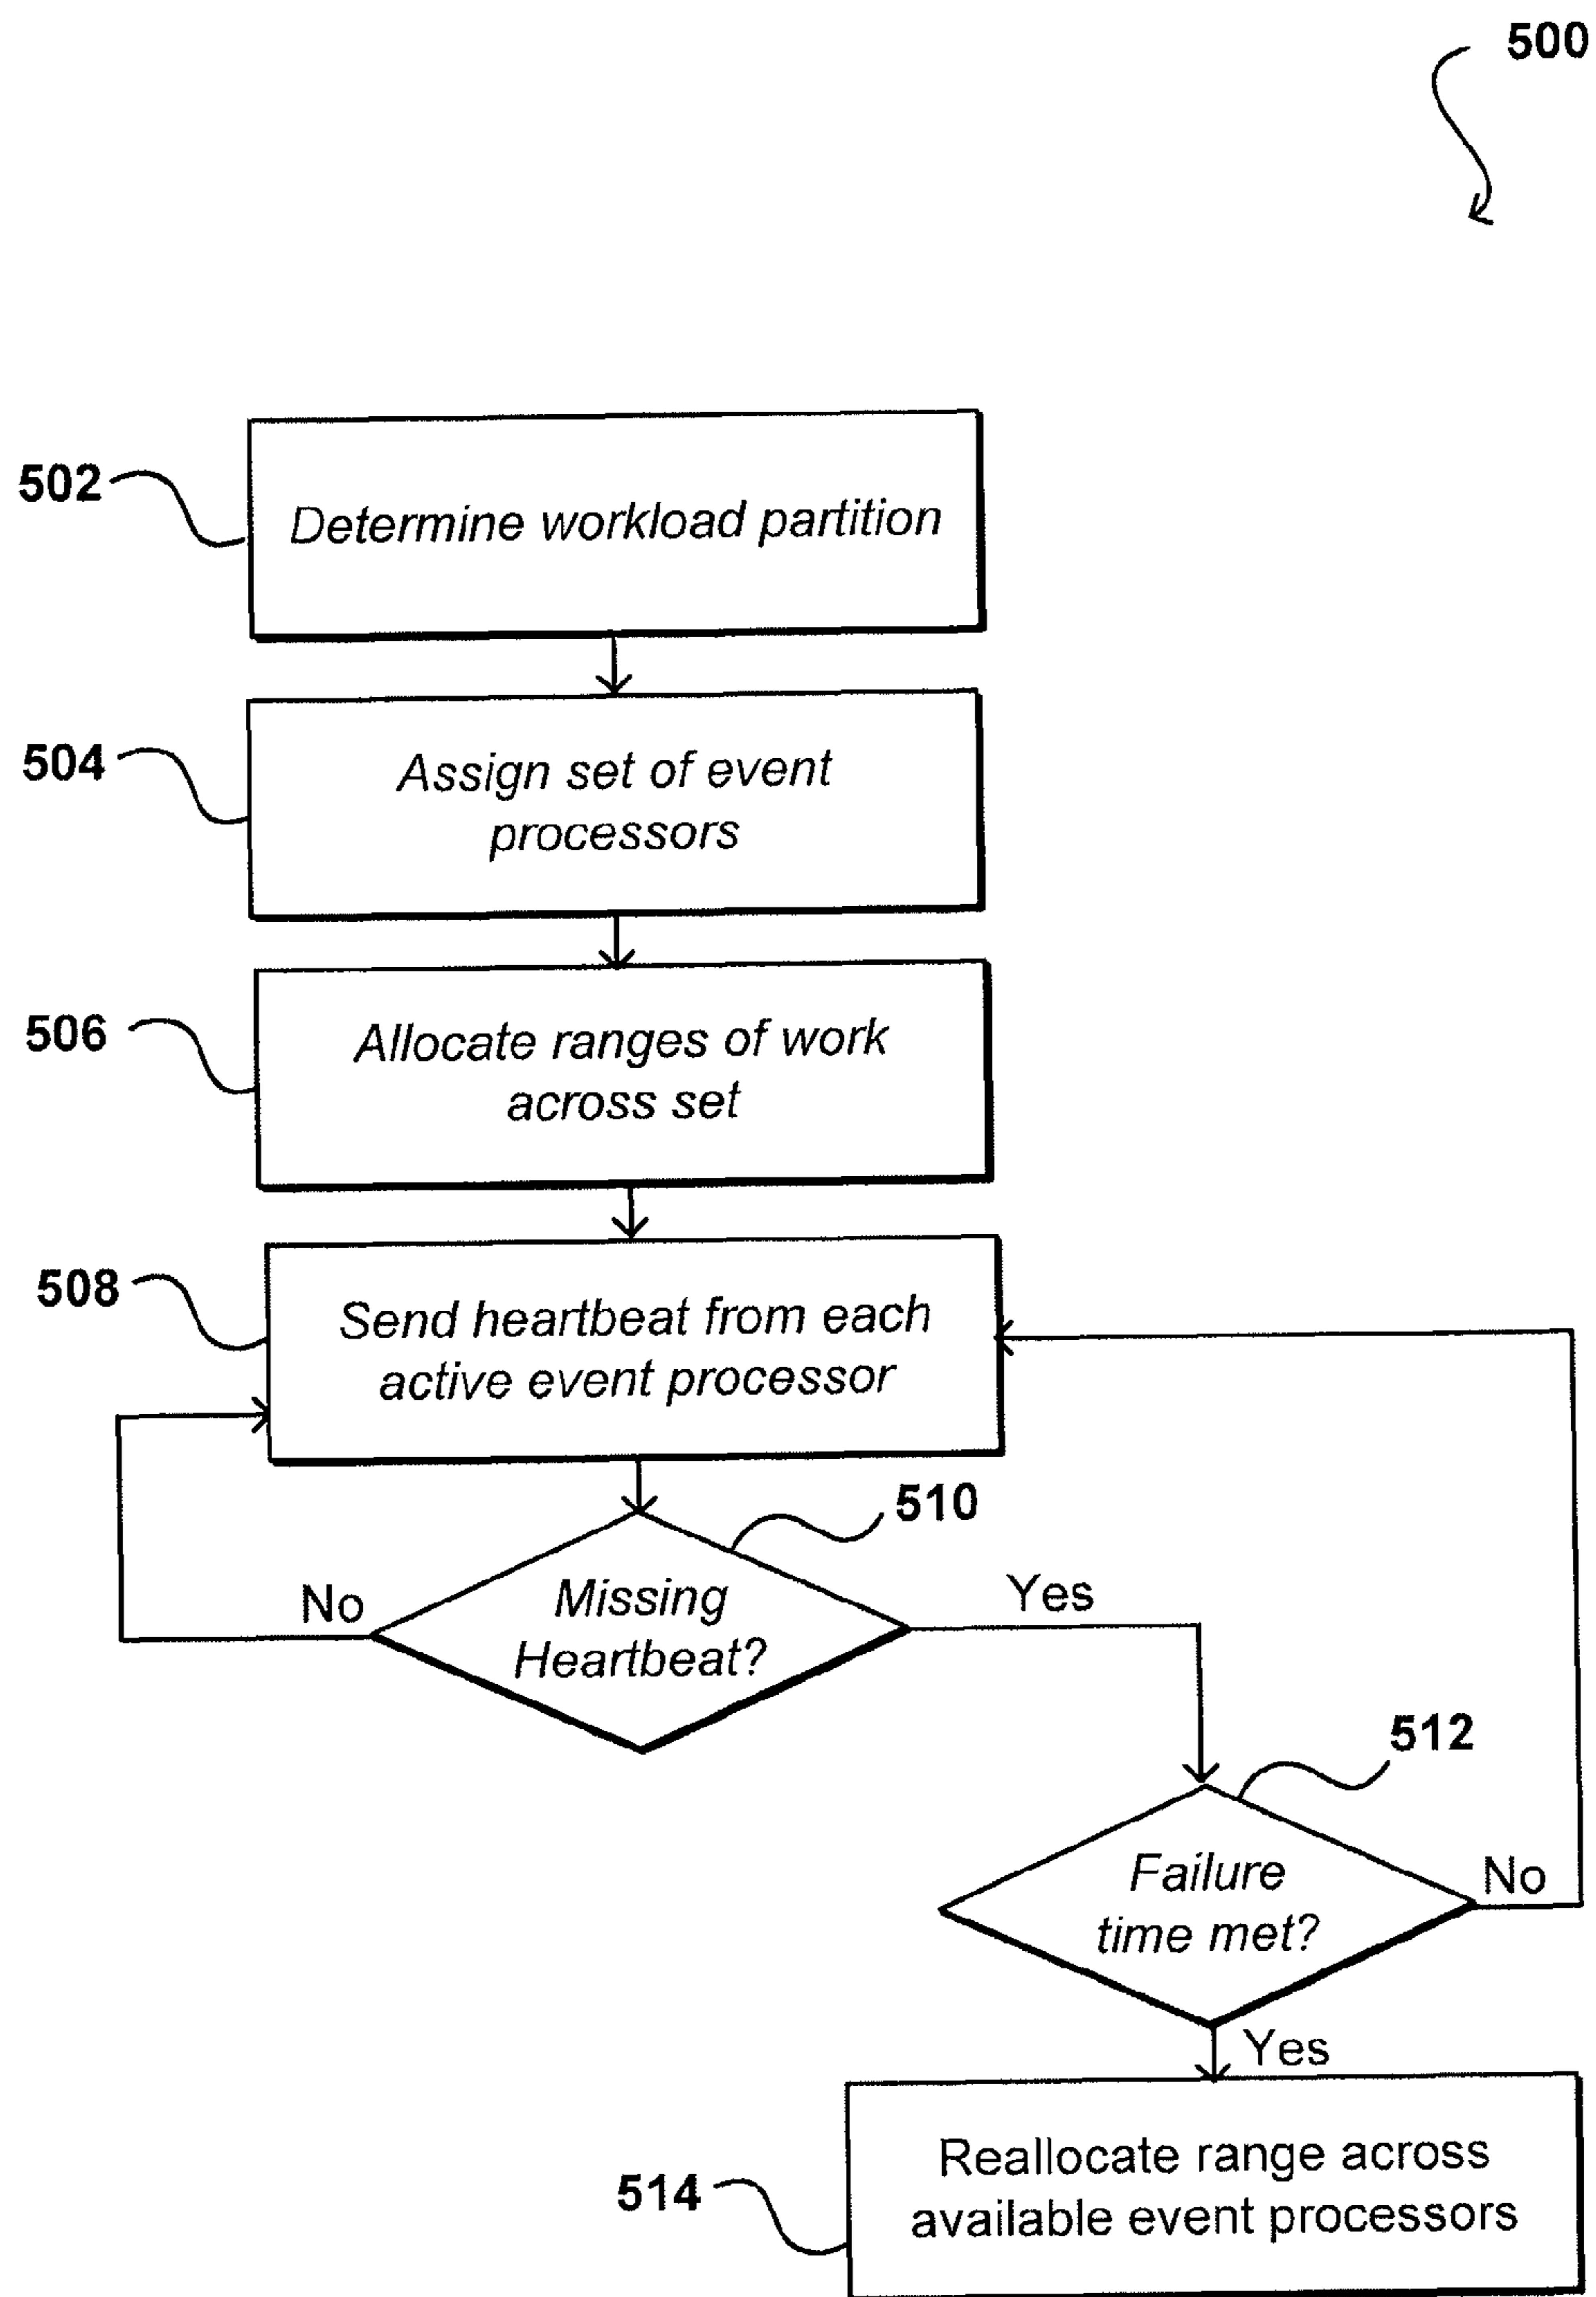


FIG. 5




FIG. 6

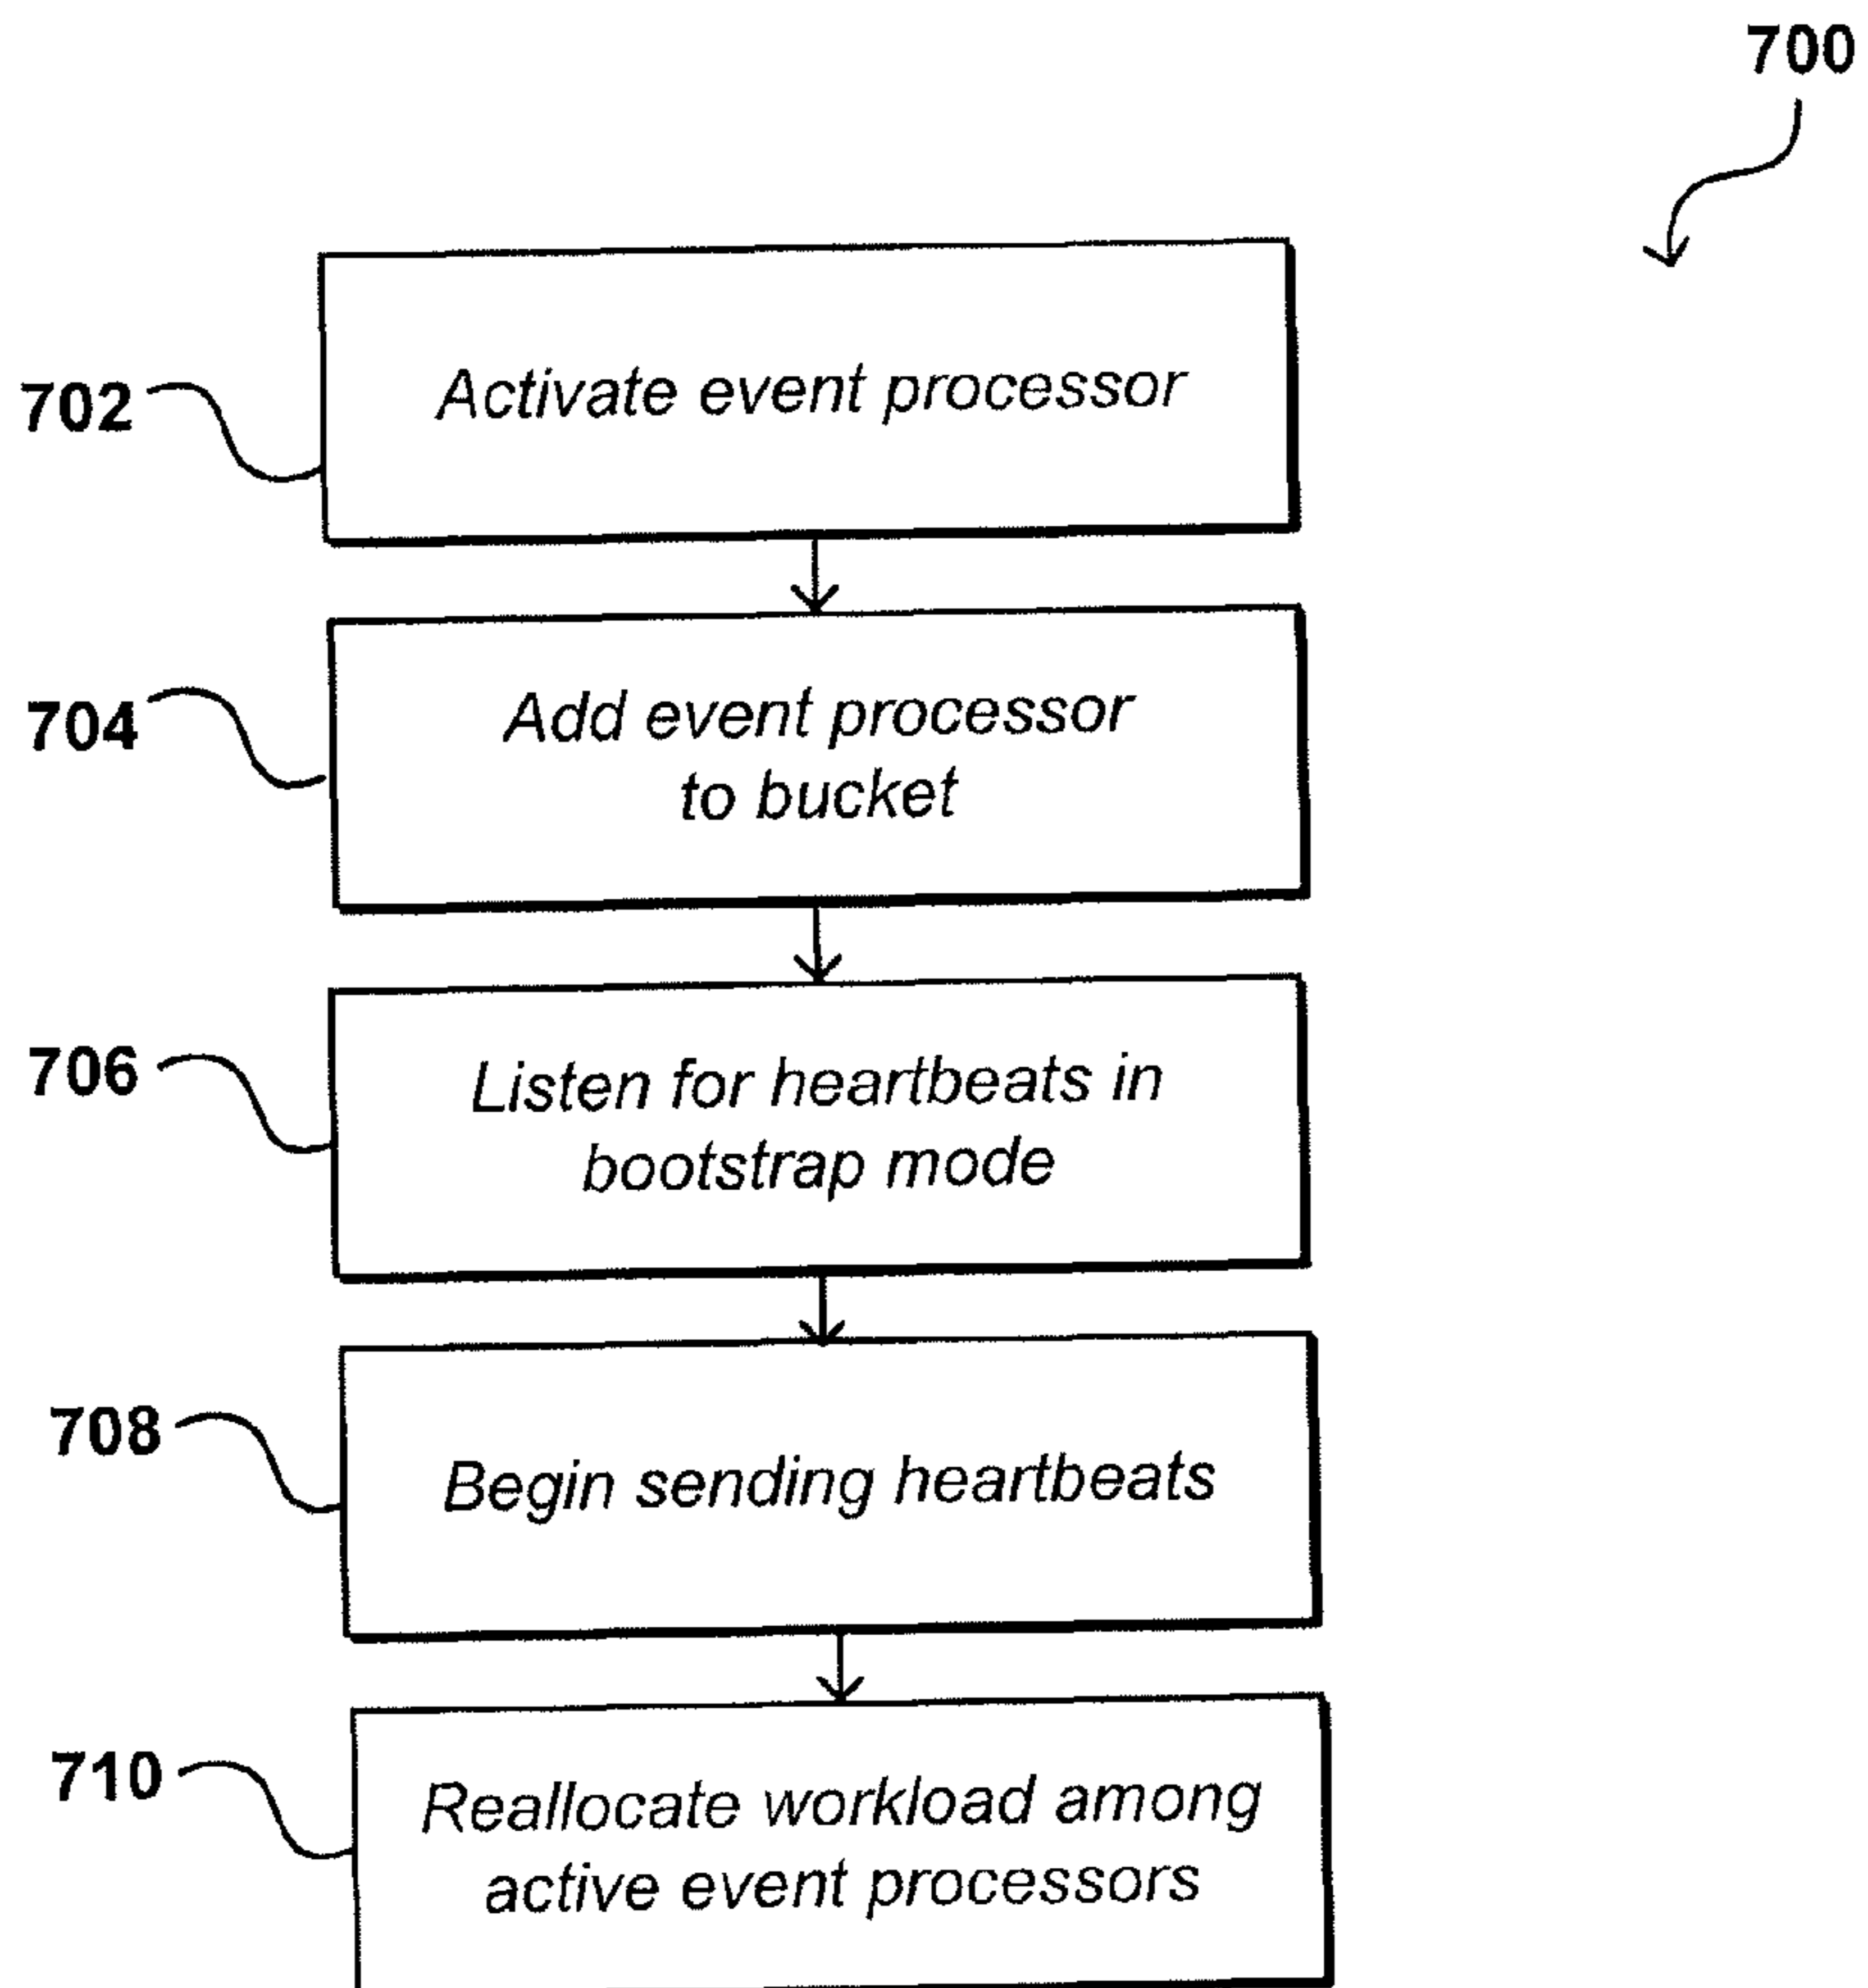


FIG. 7

1**MONITORING AND AUTOMATED
RECOVERY OF DATA INSTANCES****CROSS-REFERENCE TO RELATED
APPLICATIONS**

This application is related to co-pending U.S. patent application Ser. No. 12/415,958, entitled "Control Service and Relational Data Management," filed concurrently herewith, which is hereby incorporated herein by reference.

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND

As an increasing number of applications and services are being made available over networks such as the Internet, an increasing number of content, application, and service providers are turning to technologies such as cloud computing. Cloud computing, in general, is an approach to providing access to electronic resources through services, such as Web services, where the hardware and/or software used to support those services is dynamically scalable to meet the needs of the services at any given time. A user or customer typically will rent, lease, or otherwise pay for access to resources through the cloud, and thus does not have to purchase and maintain the hardware and/or software to provide access to these resources.

A variety of data stores and instances can be provided to a customer for use with various applications or for other such purposes. In order to prevent outages or the unavailability of customer data, it can be desirable necessary to monitor these instances for problems. This typically involves metrics being captured and logged for each instance, which allow an operator such as a database administrator (DBA) to determine the existence of potential problems, determine the type of problem, and perform an appropriate recovery. Such an approach is expensive and time consuming, however, and often results in data being unavailable for a significant amount of time while the DBA determines and corrects the problems with the data storage.

BRIEF DESCRIPTION OF THE DRAWINGS

Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

FIG. 1 illustrates an environment in which various embodiments can be implemented;

FIG. 2 illustrates an example separation of a control plane and a data plane that can be used in accordance with various embodiments;

FIG. 3 illustrates components of a control plane that can be used for monitoring and auto-recovery of elements in the data plane in accordance with various embodiments;

FIG. 4 illustrates an example process for recovering from a problem in the data plane in accordance with one embodiment;

2

FIG. 5 illustrates an example process for reallocating work in the event of an event processor failure in accordance with one embodiment;

FIG. 6 illustrates a shifting of monitored identifiers among available event processors that can be used in accordance with various embodiments; and

FIG. 7 illustrates an example process for reallocating work when an event processor comes into a bucket in accordance with one embodiment.

DETAILED DESCRIPTION

Systems and methods in accordance with various embodiments of the present disclosure may overcome one or more of the aforementioned and other deficiencies experienced in conventional approaches to managing aspects of data storage in an electronic environment. In particular, various embodiments provide a separate control environment, or control plane, that can be used to control aspects of a data environment, or data plane. The functionality of a control plane can be provided as a set of Web services, enabling the control plane to act as a virtual database administrator (DBA). A user or customer can submit a request to the control plane through an externally-visible application programming interface (API), for example, which can be analyzed to determine actions to be performed in the data plane, such as actions that create, delete, modify, expand, or otherwise modify a data store or data storage instance. Once provisioned, a user can have native access to the data instance(s) in the data plane, and can simply point existing applications (such as MySQL applications) to the DNS (domain name service) address or other location information for the particular instance. There is no restriction or modification of query models or other such functionality, as a user can continue to use applications built on MySQL, Oracle, or other such database technology.

A monitoring component of the control plane can be provided that can monitor the health or status of components in the data plane, and can automatically determine actions to be taken in the data plane. In one embodiment, a set of event processors is provided that distributes responsibility for monitoring various data instances among the event processors. The event processors can send heartbeat messages among the set, such that when a heartbeat message is not received from an event processor within a determined period of time, the other event processors in the set can automatically redistribute responsibility for monitoring among the remaining event processors.

Each active event processor can periodically check the status of each monitored component of the data plane, such as by sending a message to a responsible host manager. If the host manager sends a standard response, the event processor determines the monitored elements to be in an expected state. If no message is received, or if a message indicating an error is received, the event processor can cause information to be stored to a job queue or other data store that will result in the kickoff of a recovery workflow. The information can be analyzed to determine the likely cause of the error, and an appropriate workflow can be generated and/or determined. State information can be passed to a component of the data plane for each task necessary to perform the recovery, such that the control plane can manage the recovery without having direct access into the data stores or other such components of the data plane. Various other types of monitoring and automated recovery are provided.

FIG. 1 illustrates an example of an environment 100 for implementing aspects in accordance with various embodiments. As will be appreciated, although a Web-based envi-

environment is used for purposes of explanation, different environments may be used, as appropriate, to implement various embodiments. The environment **100** shown includes both a testing or development portion (or side) and a production portion. The production portion includes an electronic client device **102**, which can include any appropriate device operable to send and receive requests, messages, or information over an appropriate network **104** and convey information back to a user of the device. Examples of such client devices include personal computers, cell phones, handheld messaging devices, laptop computers, set-top boxes, personal data assistants, electronic book readers, and the like. The network can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network, or any other such network or combination thereof. Components used for such a system can depend at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed herein in detail. Communication over the network can be enabled by wired or wireless connections, and combinations thereof. In this example, the network includes the Internet, as the environment includes a Web server **106** for receiving requests and serving content in response thereto, although for other networks an alternative device serving a similar purpose could be used as would be apparent to one of ordinary skill in the art.

The illustrative environment includes at least one application server **108** and a data store **110**. It should be understood that there can be several application servers, layers, or other elements, processes, or components, which may be chained or otherwise configured, which can interact to perform tasks such as obtaining data from an appropriate data store. As used herein the term “data store” refers to any device or combination of devices capable of storing, accessing, and retrieving data, which may include any combination and number of data servers, databases, data storage devices, and data storage media, in any standard, distributed, or clustered environment. The application server can include any appropriate hardware and software for integrating with the data store as needed to execute aspects of one or more applications for the client device, handling a majority of the data access and business logic for an application. The application server provides access control services in cooperation with the data store, and is able to generate content such as text, graphics, audio, and/or video to be transferred to the user, which may be served to the user by the Web server in the form of HTML, XML, or another appropriate structured language in this example. The handling of all requests and responses, as well as the delivery of content between the client device **102** and the application server **108**, can be handled by the Web server. It should be understood that the Web and application servers are not required and are merely example components, as structured code discussed herein can be executed on any appropriate device or host machine as discussed elsewhere herein. Further, the environment can be architected in such a way that a test automation framework can be provided as a service to which a user or application can subscribe. A test automation framework can be provided as an implementation of any of the various testing patterns discussed herein, although various other implementations can be used as well, as discussed or suggested herein.

The environment also includes a development and/or testing side, which includes a user device **118** allowing a user such as a developer, data administrator, or tester to access the system. The user device **118** can be any appropriate device or machine, such as is described above with respect to the client device **102**. The environment also includes a development

server **120**, which functions similar to the application server **108** but typically runs code during development and testing before the code is deployed and executed on the production side and is accessible to outside users, for example. In some embodiments, an application server can function as a development server, and separate production and testing storage may not be used.

The data store **110** can include several separate data tables, databases, or other data storage mechanisms and media for storing data relating to a particular aspect. For example, the data store illustrated includes mechanisms for storing production data **112** and user information **116**, which can be used to serve content for the production side. The data store also is shown to include a mechanism for storing testing data **114**, which can be used with the user information for the testing side. It should be understood that there can be many other aspects that may need to be stored in the data store, such as for page image information and access right information, which can be stored in any of the above listed mechanisms as appropriate or in additional mechanisms in the data store **110**. The data store **110** is operable, through logic associated therewith, to receive instructions from the application server **108** or development server **120**, and obtain, update, or otherwise process data in response thereto. In one example, a user might submit a search request for a certain type of item. In this case, the data store might access the user information to verify the identity of the user, and can access the catalog detail information to obtain information about items of that type. The information then can be returned to the user, such as in a results listing on a Web page that the user is able to view via a browser on the user device **102**. Information for a particular item of interest can be viewed in a dedicated page or window of the browser.

Each server typically will include an operating system that provides executable program instructions for the general administration and operation of that server, and typically will include a computer-readable medium storing instructions that, when executed by a processor of the server, allow the server to perform its intended functions. Suitable implementations for the operating system and general functionality of the servers are known or commercially available, and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein.

The environment in one embodiment is a distributed computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components than are illustrated in FIG. 1. Thus, the depiction of the system **100** in FIG. 1 should be taken as being illustrative in nature, and not limiting to the scope of the disclosure.

An environment such as that illustrated in FIG. 1 can be useful for a provider such as an electronic marketplace, wherein multiple hosts might be used to perform tasks such as serving content, authenticating users, performing payment transactions, or performing any of a number of other such tasks. Some of these hosts may be configured to offer the same functionality, while other servers might be configured to perform at least some different functions. The electronic environment in such cases might include additional components and/or other arrangements, such as those illustrated in the configuration **200** of FIG. 2, discussed in detail below.

Systems and methods in accordance with one embodiment provide a relational database service (“RDS”) that enables developers, customers, or other authorized users to easily and

cost-effectively obtain and configure relational databases so that users can perform tasks such as storing, processing, and querying relational data sets in a cloud. While this example is discussed with respect to the Internet, Web services, and Internet-based technology, it should be understood that aspects of the various embodiments can be used with any appropriate services available or offered over a network in an electronic environment. Further, while the service is referred to herein as a “relational database service,” it should be understood that such a service can be used with any appropriate type of data repository or data storage in an electronic environment. An RDS in this example includes at least one Web service that enables users or customers to easily manage relational data sets without worrying about the administrative complexities of deployment, upgrades, patch management, backups, replication, failover, capacity management, scaling, and other such aspects of data management. Developers are thus freed to develop sophisticated cloud applications without worrying about the complexities of managing the database infrastructure.

An RDS in one embodiment provides a separate “control plane” that includes components (e.g., hardware and software) useful for managing aspects of the data storage. In one embodiment, a set of data management application programming interfaces (APIs) or other such interfaces are provided that allow a user or customer to make calls into the RDS to perform certain tasks relating to the data storage. The user still can use the direct interfaces or APIs to communicate with the data repositories, however, and can use the RDS-specific APIs of the control plane only when necessary to manage the data storage or perform a similar task.

FIG. 2 illustrates an example of an RDS implementation **200** that can be used in accordance with one embodiment. In this example, a computing device **202** for an end user is shown to be able to make calls through a network **206** into a control plane **208** to perform a task such as to provision a data repository of the data plane **210**. The user or an application **204** can access the provisioned repository directly through an interface of a data plane **210**. While an end user computing device and application are used for purposes of explanation, it should be understood that any appropriate user, application, service, device, component, or resource can access the interface(s) of the control plane and/or data plane as appropriate in the various embodiments. Further, while the components are separated into control and data “planes,” it should be understood that this can refer to an actual or virtual separation of at least some resources (e.g., hardware and/or software) used to provide the respective functionality.

The control plane **208** in this example is essentially a virtual layer of hardware and software components that handles control and management actions, such as provisioning, scaling, replication, etc. The control plane in this embodiment includes a Web services layer **212**, or tier, which can include at least one Web server, for example, along with computer-executable software, application servers, or other such components. The Web services layer also can include a set of APIs **232** (or other such interfaces) for receiving Web services calls or requests from across the network **206**, which the Web services layer can parse or otherwise analyze to determine the steps or actions needed to act on or process the call. For example, a Web service call might be received that includes a request to create a data repository. In this example, the Web services layer can parse the request to determine the type of data repository to be created, the storage volume requested, the type of hardware requested (if any), or other such aspects. Information for the request can be written to an

administration (“Admin”) data store **222**, or other appropriate storage location or job queue, for subsequent processing.

A Web service layer in one embodiment includes a scalable set of customer-facing servers that can provide the various control plane APIs and return the appropriate responses based on the API specifications. The Web service layer also can include at least one API service layer that in one embodiment consists of stateless, replicated servers which process the customer APIs. The Web service layer can be responsible for Web service front end features such as authenticating customers based on credentials, authorizing the customer, throttling customer requests to the API servers, validating user input, and marshalling or unmarshalling requests and responses. The API layer also can be responsible for reading and writing database configuration data to/from the administration data store, in response to the API calls. In many embodiments, the Web services layer will be the only externally visible component, or the only component that is visible to, and accessible by, customers of the control service. The servers of the Web services layer can be stateless and scaled horizontally as known in the art. API servers, as well as the persistent data store, can be spread across multiple data centers in a region, for example, such that the servers are resilient to single data center failures.

The control plane in this embodiment includes what is referred to herein as a “sweeper” component **214**. A sweeper component can be any appropriate component operable to poll various components of the control plane or otherwise determine any tasks to be executed in response to an outstanding request. In this example, the Web services layer might place instructions or information for the “create database” request in the admin data store **222**, or a similar job queue, and the sweeper can periodically check the admin data store for outstanding jobs. Various other approaches can be used as would be apparent to one of ordinary skill in the art, such as the Web services layer sending a notification to a sweeper that a job exists. The sweeper component can pick up the “create database” request, and using information for the request can send a request, call, or other such command to a workflow component **216** operable to instantiate at least one workflow for the request. The workflow in one embodiment is generated and maintained using a workflow service as is discussed elsewhere herein. A workflow in general is a sequence of tasks that should be executed to perform a specific job. The workflow is not the actual work, but an abstraction of the work that controls the flow of information and execution of the work. A workflow also can be thought of as a state machine, which can manage and return the state of a process at any time during execution. A workflow component (or system of components) in one embodiment is operable to manage and/or perform the hosting and executing of workflows for tasks such as: repository creation, modification, and deletion; recovery and backup; security group creation, deletion, and modification; user credentials management; and key rotation and credential management. Such workflows can be implemented on top of a workflow service, as discussed elsewhere herein. The workflow component also can manage differences between workflow steps used for different database engines, such as MySQL, as the underlying workflow service does not necessarily change.

In this example, a workflow can be instantiated using a workflow template for creating a database and applying information extracted from the original request. For example, if the request is for a MySQL® Relational Database Management System (RDBMS) instance, as opposed to an Oracle® RDBMS or other such instance, then a specific task will be added to the workflow that is directed toward MySQL

instances. The workflow component also can select specific tasks related to the amount of storage requested, any specific hardware requirements, or other such tasks. These tasks can be added to the workflow in an order of execution useful for the overall job. While some tasks can be performed in parallel, other tasks rely on previous tasks to be completed first. The workflow component or service can include this information in the workflow, and the tasks can be executed and information passed as needed.

An example “create database” workflow for a customer might include tasks such as provisioning a data store instance, allocating a volume of off-instance persistent storage, attaching the persistent storage volume to the data store instance, then allocating and attaching a DNS (domain name system) address or other address, port, interface, or identifier which the customer can use to access or otherwise connect to the data instance. In this example, a user is provided with the DNS address and port to be used to access the instance. The workflow also can include tasks to download and install any binaries or other information used for the specific data storage technology (e.g., MySQL). The workflow component can manage the execution of these and any related tasks, or any other appropriate combination of such tasks, and can generate a response to the request indicating the creation of a “database” in response to the “create database” request, which actually corresponds to a data store instance in the data plane **210**, and provide the DNS address to be used to access the instance. A user then can access the data store instance directly using the DNS address and port, without having to access or go through the control plane **208**. Various other workflow templates can be used to perform similar jobs, such as deleting, creating, or modifying one of more data store instances, such as to increase storage. In some embodiments, the workflow information is written to storage, and at least one separate execution component (not shown) pulls or otherwise accesses or receives tasks to be executed based upon the workflow information. For example, there might be a dedicated provisioning component that executes provisioning tasks, and this component might not be called by the workflow component, but can monitor a task queue or can receive information for a provisioning task in any of a number of related ways as should be apparent.

As mentioned, various embodiments can take advantage of a workflow service that can receive requests or calls for a current state of a process or task, such as the provisioning of a repository, and can return the current state of the process. The workflow component and/or workflow service do not make the actual calls or requests to perform each task, but instead manage the state and configuration information for the workflow that enables the components of the control plane to determine the next task to be performed, and any information needed for that task, then generate the appropriate call(s) into the data plane including that state information, whereby a component of the data plane can make the call to perform the task. Workflows and tasks can be scheduled in parallel in order to increase throughput and maximize processing resources. As discussed, the actual performing of the tasks will occur in the data plane, but the tasks will originate from the control plane. For example, the workflow component can communicate with a host manager, which can make calls into the data store. Thus, for a given task a call could be made to the workflow service passing certain parameters, whereby the workflow service generates the sequence of tasks for the workflow and provides the current state, such that a task for the present state can be performed. After the task is performed (or otherwise resolved or concluded), a component such as the host manager can reply to the service, which can then

provide information about the next state in the workflow, such that the next task can be performed. Each time one of the tasks for the workflow is performed, the service can provide a new task to be performed until the workflow is completed. Further, multiple threads can be running in parallel for different workflows to accelerate the processing of the workflow.

The control plane **208** in this embodiment also includes at least one monitoring component **218**. When a data instance is created in the data plane, information for the instance can be written to a data store in the control plane, such as a monitoring data store **220**. It should be understood that the monitoring data store can be a separate data store, or can be a portion of another data store such as a distinct set of tables in an Admin data store **222**, or other appropriate repository. A monitoring component can access the information in the monitoring data store to determine active instances **234** in the data plane **210**. A monitoring component also can perform other tasks, such as collecting log and/or event information from multiple components of the control plane and/or data plane, such as the Web service layer, workflow component, sweeper component, and various host managers. Using such event information, the monitoring component can expose customer-visible events, for purposes such as implementing customer-facing APIs. A monitoring component can constantly monitor the health of all the running repositories and/or instances for the control plane, detect the failure of any of these instances, and initiate the appropriate recovery process(es).

Each instance **234** in the data plane can include at least one data store **226** and a host manager component **228** for the machine providing access to the data store. A host manager in one embodiment is an application or software agent executing on an instance and/or application server, such as a Tomcat or Java application server, programmed to manage tasks such as software deployment and data store operations, as well as monitoring a state of the data store and/or the respective instance. A host manager in one embodiment listens on a port that can only be reached from the internal system components, and is not available to customers or other outside entities. In some embodiments, the host manager cannot initiate any calls into the control plane layer. A host manager can be responsible for managing and/or performing tasks such as setting up the instances for a new repository, including setting up logical volumes and file systems, installing database binaries and seeds, and starting or stopping the repository. A host manager can monitor the health of the data store, as well as monitoring the data store for error conditions such as I/O errors or data storage errors, and can restart the data store if necessary. A host manager also perform and/or manage the installation of software patches and upgrades for the data store and/or operating system. A host manager also can collect relevant metrics, such as may relate to CPU, memory, and I/O usage.

The monitoring component can communicate periodically with each host manager **228** for monitored instances **234**, such as by sending a specific request or by monitoring heartbeats from the host managers, to determine a status of each host. In one embodiment, the monitoring component includes a set of event processors (or monitoring servers) configured to issue commands to each host manager, such as to get the status of a particular host and/or instance. If a response is not received after a specified number of retries, then the monitoring component can determine that there is a problem and can store information in the Admin data store **222** or another such job queue to perform an action for the instance, such as to verify the problem and re-provision the instance if necessary. The sweeper can access this information and kick off a recovery workflow for the instance to attempt to automatically

recover from the failure. The host manager **228** can act as a proxy for the monitoring and other components of the control plane, performing tasks for the instances on behalf of the control plane components. Occasionally, a problem will occur with one of the instances, such as the corresponding host, instance, or volume crashing, rebooting, restarting, etc., which cannot be solved automatically. In one embodiment, there is a logging component (not shown) that can log these and other customer visibility events. The logging component can include an API or other such interface such that if an instance is unavailable for a period of time, a customer can call an appropriate “events” or similar API to get the information regarding the event. In some cases, a request may be left pending when an instance fails. Since the control plane in this embodiment is separate from the data plane, the control plane never receives the data request and thus cannot queue the request for subsequent submission (although in some embodiments this information could be forwarded to the control plane). Thus, the control plane in this embodiment provides information to the user regarding the failure so the user can handle the request as necessary.

As discussed, once an instance is provisioned and a user is provided with a DNS address or other address or location, the user can send requests “directly” to the data plane **210** through the network using a Java Database Connectivity (JDBC) or other such client to directly interact with that instance **234**. In one embodiment, the data plane takes the form of (or at least includes or is part of) a computing cloud environment, or a set of Web services and resources that provides data storage and access across a “cloud” or dynamic network of hardware and/or software components. A DNS address is beneficial in such a dynamic cloud environment, as instance or availability failures, for example, can be masked by programmatically remapping a DNS address to any appropriate replacement instance for a use. A request received from a user **202** or application **204**, for example, can be directed to a network address translation (NAT) router **224**, or other appropriate component, which can direct the request to the actual instance **234** or host corresponding to the DNS of the request. As discussed, such an approach allows for instances to be dynamically moved, updated, replicated, etc., without requiring the user or application to change the DNS or other address used to access the instance. As discussed, each instance **234** can include a host manager **228** and a data store **226**, and can have at least one backup instance or copy in persistent storage **230**. Using such an approach, once the instance has been configured through the control plane, a user, application, service, or component can interact with the instance directly through requests to the data plane, without having to access the control plane **232**. For example, the user can directly issue structured query language (SQL) or other such commands relating to the data in the instance through the DNS address. The user would only have to access the control plane if the user wants to perform a task such as expanding the storage capacity of an instance. In at least one embodiment, the functionality of the control plane **208** can be offered as at least one service by a provider that may or may not be related to a provider of the data plane **210**, but may simply be a third-party service that can be used to provision and manage data instances in the data plane, and can also monitor and ensure availability of those instances in a separate data plane **210**.

As discussed, one advantage to use of a control plane is that the control plane can function as a virtual database administrator (DBA) and avoid the need for a human DBA to perform tasks such as performing recovery actions in the event of a component failure. In conventional systems, metrics or other

such information are collected and a DBA is tasked with analyzing the information. Failure of a data instance in the cloud can lead to a data store outage during the time needed for recovery. In order to recover from an instance failure, the failure of the instance first needs to be detected. Once the failure is detected, the type of failure needs to be determined, and an appropriate recovery action taken. Conventional approaches relying on a DBA to perform these actions are expensive and time-consuming, and can result in significant unavailability of customer data during the detection and recovery process.

FIG. **3** illustrates an example of a configuration **300** that can be used for purposes such as monitoring and automated recovery in accordance with one embodiment. Although reference numbers are carried over between figures for purposes of simplicity and clarity, it should be understood that these merely represent similar components that can be used for various embodiments, and should not be interpreted as requiring components from various other embodiments or as merely showing different views of a single embodiment. Further, fewer or additional components can be used in various embodiments, and the presence or lack of a component in a given figure should not be interpreted as that component being required or not useful in a given embodiment unless otherwise specifically stated. Variations between the embodiments and figures should be apparent to one of ordinary skill in light of the present disclosure.

As illustrated in the figure, a monitoring component (or service) **218** of the control plane can comprise a series of processing nodes **302**, referred to herein as event processors. In one embodiment, the event processors comprise a fleet of monitoring servers operable to monitor aspects of the data plane. Each event processor can be configured to communicate with a specified set or range of data stores **226** and/or data instances **234** through the associated host manager **228**. As discussed, each data store and host manager can exist on a node or machine of the data plane **210**, or data environment. Each of the event processors can communicate with the allocated host managers using any appropriate communication technique to obtain a current status from each host, such as by pinging each host manager using a secure (e.g., HTTPS) request, such as a “getStatus” request. In response to the request, each host manager can send a response including information such as whether there is a problem with, or detected by, the host manager **228**, as well as any relevant metrics, parameter values, or diagnostic information that is determined to be relevant. In certain embodiments, the amount and type of information returned by a host manager can vary based upon a state of the host manager. For example, if there are no errors detected then the host manager might send a standard set of specified metrics to be logged or otherwise processed. If a problem is detected, for example, then a different set of information might be included, such as information indicating the type of problem as well as diagnostic or other information relevant to that type of problem. Various algorithms can be provided to the host managers for making such determinations. Upon receiving the information from the host managers, the event processors can analyze the information, as necessary, and store the information in a monitoring data store **220** or other such location. The event processors can also store any log information, discussed elsewhere herein, in the monitoring data store. As illustrated in this example, the monitoring data store **220** can be a single logical data store, but can be partitioned across many data instances **304**.

There can be many advantages to using multiple event processors **302** as part of the monitoring component **218**. One

such advantage is that, for a large number of data instances **234** in the data plane, a single event processor may not have enough capacity to monitor each instance concurrently. Utilizing multiple event processors allows the monitoring work to be distributed across several event processors. Further, using multiple event processors allows for existing event processors to take on the work of another event processor in the event of a failure or other such problem. If a data instance was only managed by a single event processor, and there was a problem with that processor making the event processor unavailable, then that data instance might not have any monitoring performed and thus could risk an outage or other such problem. By spreading the monitoring across a set of event processors, and allowing the range of monitoring by each event processor to update dynamically, the control plane can ensure that each instance in the data plane is monitored at substantially any time, even in the event of a failure of one or more of the event processors.

In one embodiment, the responsibility of each event processor is determined by taking the number of instances to be monitored at any given time and apportioning the number of instances across the number of event processors. For example, if there are 25,000 instances to be monitored in the data plane, and there are five event processors running in the control plane, then each event processor can be given responsibility for monitoring approximately 5,000 of the data instances. If each instance is given an identifier, for example, then each event processor can be given a range of identifiers (such as the first 5,000 identifiers, second 5,000 identifiers, etc.) to make it easier to adjust responsibility for each event processor, rather than having to manage mapping information for each of the 25,000 instances. The example in the figure shows the range of responsibilities for each of the event processors in such an example.

At an appropriate interval, such as once a minute, each event processor **302** can send a request to each host manager **228** being monitored by that event processor. An event processor in one embodiment is a Java application running within a Tomcat container of the control plane that regularly polls the host managers for data instances in the data plane. The event processor can poll a host manager in one embodiment by making a `getStatus()` or similar call (e.g., over SSL) using the DNS name and host manager port. In some embodiments a data instance being monitored is uniquely identified by a combination of a customer data store identifier, a data store identifier, and an instance identifier. Using such an approach, the states of the old and new instances can be distinguished when moving a data instance to another instance in the cloud. The event processor can determine the state of the data instance based upon the response from of the host manager. A data instance in one embodiment can be in one of at least the following example states: “OK” (the data instance is running properly), “incommunicado” (the data instance is in a suspect state of failure), or “dead” (the data instance is unreachable and does not respond to requests for status).

In most cases, the host manager will return a response indicating that the host manager, associated instance, etc., is running as expected, and the event processor can update information in the monitoring data store **220**. An event processor can consider a data instance to be in an “OK” or similar state in one embodiment when the host manager returns an appropriate response, such as an HTTP response code “200” (a standard response code for successful HTTP requests). If a response is not received from a host manager, or if the response is a timed-out response (such as HTTP code “500”, or any other “5xx” error response codes), the event processor

can resend the `getStatus` request, and can place the database instance in an “incommunicado” or similar state. If the host has been in the “incommunicado” state for more than a predetermined number of status pings, or other such requests, then the data instance can be declared to be in a “dead” or similar state. If the host comes back online with a “200” response (or similar) code within the predetermined number of status pings, the host or instance can be moved to an “OK” state. The predetermined number of checks before moving a host state from “incommunicado” to “dead” or “OK” used, at least in part, is to avoid false positives due to intermittent network errors, temporarily overloaded event processors, temporarily overloaded host managers, or other such temporary errors that do not actually result in a data instance being unavailable other otherwise requiring recovery. In one embodiment, a state of “incommunicado” is not persisted, as the state can easily be determined by another event processor.

If a reply is not received after the predetermined number of status requests, or the state is otherwise moved to a “dead” or similar state, the event processor enters information regarding the problem state into the Admin data store **222** (or other such job queue as discussed above) indicating that there is a suspect state with respect to the unresponsive host manager. As discussed above, a sweeper **214** component of the control plane can periodically check the Admin data store for information, and when the sweeper detects the information for the suspect or problem state, an appropriate recovery workflow can be started. For example, the sweeper can pass information to the workflow component **216** that causes an appropriate workflow to be generated, such as a workflow to handle a data instance being unavailable, a workflow to handle errors reported by a host manager, or any of a number of other such situations. The workflow manager can generate the appropriate workflow, pass state information, and handle various other aspects as discussed elsewhere herein.

One advantage to storing recovery information in the Admin data store is that such an approach allows for recovery even in the event of a failure of the monitoring system. It can be desirable to enable recovery actions independent of the availability of the monitoring data store. It can be acceptable to use the Admin data store, as in this embodiment any type of recovery, including generating a workflow, etc., requires the Admin data store (or other such job queue) to be active and available. It can thus be desirable to avoid placing another dependency on the recovery, and instead having a single place of availability.

FIG. 4 illustrates an example of a process **400** for detecting and recovering from a failure in the data environment in accordance with one embodiment, using components and approaches such as those discussed above with respect to FIG. 3. In this example, at least one workload partition is determined for the data plane **402**. Depending at least in part upon the number of data stores, instances, host managers, and other such components to be monitored, the overall workload may be partitioned into any of a number of separate partitions. A set of event processors can be assigned to each workload partition **404**, and each event processor in the set is allocated a respective portion of the work for the assigned partition **406**. At the appropriate intervals, each event processor sends status requests to each host manager (or other such component) being monitored by that event processor **408**. The event processors wait to receive responses from those host managers **410**. A determination is made as to whether responses are received from each host manager within a predetermined amount of time **412**, or whether a timed-out response is received indicating that a host manager did not generate a response, and if so a determination can be made as to whether

13

that host manager has failed to respond for a threshold number of requests **414**. If the threshold number of requests has not been sent for which no replies have been received, the request can be resent **416**. If the threshold number of requests has been sent without response, information relating to the respective host manager can be written to the Admin data store **418** or another appropriate data store or job queue.

If all the responses were received, a determination is made as to whether any of the responses indicated the presence of a potential problem **420**. If not, the control plane can proceed and the event processors will send requests again at the next interval. If there are problems, this information also can be written to the Admin data store **418**. As discussed, a sweeper or other such component can detect the information in the data store and cause an appropriate recovery workflow to be kicked off **422**. The appropriate state information can be passed to the data plane in order to perform the recovery **424**, after which the event processor responsible for the component in the data plane can return to sending status requests at the appropriate intervals. If the recovery is not successful, other actions can be taken, such as notifying a DBA or other such operator.

The appropriate action to be taken can depend, at least in part, upon the type of problem detected. For example, failure of a data instance or host manager can result in the generation and execution of a recovery workflow. If the problem is a component such as a host machine rebooting, for example, then no recovery workflow may be needed. In some cases, such as for the failure of an entire data center, there may be a number of concurrent problems detected. In one embodiment, up to a threshold number of errors can occur that are handled using recovery or other such workflows. If the number exceeds that threshold, a system in accordance with one embodiment can page or otherwise notify an operator to analyze the problems and determine an appropriate course of action. In other embodiments, a staged recovery can be initiated wherein a limited number of recovery workflows are executed at any given time to avoid problems such as overloading the system.

As discussed, each partition of the monitoring workload can be covered by a set of event processors. Covering a single partition of the workload with a set of event processors enables the redistributing of the monitoring load across the remaining event processors in the event that one of the event processors fail or experiences any of a variety of other such problems. In one embodiment, each group of event processors is contained in a bucket or other such partition. Each event processor in a bucket is responsible for handling a range of instances in a single data plane, or grouping of instances in that plane. A failure detection process can be used to ensure that if a failure occurs, the other event processors in that bucket take over responsibility for the instances handled by the failed event processor. The monitoring data store in at least one embodiment holds the list of current active data instances to be monitored by the set of event processors in a bucket, as well as the information that the event processors collect for various customer-related events. As the number of monitored instances increases, it can be necessary to scale beyond a single monitoring data store. Thus, each table in the monitoring data store can be partitioned, including the db_poll_list.

14

In one embodiment, the event processors are deployed with a partition table of the following example format:

Partition Id	Hash Range
P0	0-10000
P1	10000-20000

This partition configuration can be deployed as a configuration file to the event processor hosts.

If a given workload partition generates a significant number of events that leaves the responsible set of event processors in a constant catch-up mode (i.e., not able to finish the assigned health checks within a certain time period), additional event processors can be added to the set responsible for that workload partition without having to repartition the data store. Using such a technique, the performance scalability can be differentiated from the data scalability issues. For example, a single partition generating so many events that the event processors cannot catch up can be distinguished from a situation where the single partition generate so many events that a single data store does not provide enough storage space.

The membership of the event processors and the partitions to which the event processors are assigned can be stored in a location such as an event processor membership configuration file. The membership configuration information can be deployed to the event processors in a group (such as in the same partition or bucket), and can have the following example format:

```
<EP identifier> <EP Host Name> <endpoint_port> <Partition Id>
```

When a single partition is covered by multiple event processors, each event processor splits the bucket ranges by sorting the event processor identifiers, such as by using a lexicographic or hash-based sorting routine, and dividing the bucket ranges uniformly. Each event processor can independently determine the appropriate range to be monitored.

FIG. 5 illustrates an example process **500** for monitoring the health of event processors in a bucket and handling the failure of one of the event processors in accordance with one embodiment. In this example, at least one workload partition is determined for the data plane **502**. Depending at least in part upon the number of data stores, instances, host managers, and other such components to be monitored, the overall workload may be partitioned into any of a number of separate partitions. A set of event processors can be assigned to each workload partition **504**, and each event processor in the set is allocated a respective portion of the work for the assigned partition **506**. At the appropriate intervals, each event processor sends a "heartbeat" message (e.g., over HTTP) to the event processors in the same set or bucket covering the same workload partition **508**. The heartbeats can be sent at any appropriate interval, such as every ten seconds. A "heartbeat" in one embodiment refers to a simple multicast message that is sent to each event processor in a bucket to inform the other event processors of the status of the event processor sending the heartbeat. The event processors can maintain a list of event processors and their status (e.g., "available" or "dead") along with the last check-in time of each event processor. If it is determined that a heartbeat is received from each event processor in the bucket **510**, the process can continue.

If, however, it is determined that an event processor in the same bucket has not responded with a heartbeat, then a determination is made as to whether the event processor has failed to send a heartbeat for a time period equal to, or greater than,

a specified heartbeat failure time (e.g., be six times the heartbeat interval) **512**. If the specified heartbeat failure time has not been reached, the process can continue. If the heartbeat failure time has been at least reached without a heartbeat from an event processor, each active event processor in the bucket can declare the non-responsive event processor to be “dead”, or in a similar state, and can reallocate the responsibility ranges and take over a portion of the monitoring workload **514**. As every active event processor in the bucket will fail to receive a heartbeat message from the failed event processor, the event processors can each expand the allocated workload by an appropriate amount to pick up the work of the “missing” event processor.

As illustrated in the example **600** of FIG. 6, if there are four event processors and 60,000 instances being monitored, then each event processor handles 15,000 instances (which can be ordered in lexicographical order or another appropriate order by identifier, etc.). If one of the event processors fails, the other three event processors can re-allocate their respective range of responsibility, such that each event processor now handles 20,000 of the instances (still being consecutively ordered according to the identifier, etc). Thus, since the instances are ordered using an ordering scheme, the event processors can adjust the range of the ordering scheme to be monitored, and do not have to map or otherwise track which “new” instances to monitor. The ranges being monitored can be stored in the monitoring data store, for example. Such an approach is also beneficial in situations where instances are added or removed, as the workload can be automatically distributed (substantially) evenly across the event processors. Heartbeating only within a particular bucket also can be more efficient and easy to maintain than a global heartbeating mechanism.

FIG. 7 illustrates an example process **700** for reallocating the work ranges across a bucket when an event processor is added to the bucket, such as may be a result of adding additional processing capacity or a result of a failed event processor recovering and again being able to handle a portion of the workload. An event processor can become active **702**, such as by an event processor host restarting or recovering, or the host simply being activated or added to a bucket. The event processor also can be added to the bucket **704**, although in cases of recovery the event processor might already be assigned to that bucket. Upon the active event processor being added to the bucket, the event manager can enter a mode such as a “bootstrap” mode for a time period (e.g., the heartbeat failure time) to receive heartbeats from the peer event processors in the bucket **706**, to obtain information about the other event processors active in the bucket and determine a time for sending heartbeats, for example. The event processor can engage a heartbeating agent to also start sending heartbeats to the other event processors in the bucket **708**. After this time, the host can move itself to an “operational” mode, where each event processor can reallocate the range of work and determines its current slice of monitoring workload based on the state of the event processors assigned to its partition **710**. One reason for leaving the event processors in “bootstrap” mode for a period of time is to ensure that the new event processor that joins (or rejoins) the event processor collective, and the remaining event processors, have sufficient time to converge on the current state of active event processors.

An approach in accordance with one embodiment also over-partitions the event processors, such as by running each event processor at 50-60% of capacity. Such an approach enables at least one or two event processors to fail in each bucket without having a significantly negative impact on performance. A failed event processor will eventually become

available again, such as where the respective host reboots. That event processor then can start exchanging heartbeats again, whereby the other event processors in the bucket can automatically detect the presence of the event processor. The allocated work can be automatically redistributed as discussed above, so that the work is relatively evenly distributed across the larger set of available event processors in the bucket.

In such a system, it can also be important to ensure that the list or set of data stores and/or instances to be monitored is automatically populated and updated over time. One approach would be to create a database list table, for example, which is a snapshot replica of the instances which can be propagated as needed. Such an approach, however, can be difficult to maintain, as well as to ensure that each appropriate component has the most recent copy. Another approach would be to have the event processors query the data plane components, and then store the information locally in the control plane. Such an approach can create a lot of messaging traffic, and can be difficult to maintain and update. An approach in accordance with one embodiment instead enables each event processor to expose an interface such as a “setStatus” or similar API. As part of a “create” or “delete” workflow, for example, a task can be added to the end of the workflow which instructs the appropriate host manager to call the event processor that is, or was, in charge of managing the instance. The host manager can thus call the “setStatus” API of the event processor to set a status of the host, any time there is a change in status as a result of a workflow (or other such action). Each time an event processor receives a call through the “setStatus” API, information can be placed in a local data store to add the new host to its set of partitions, remove the host, etc. Information for the host also can be written to the monitoring data store or another appropriate persistent location.

In one embodiment, an authoritative list of current active data instances resides in the Admin data store. An active list of data instances to be monitored resides in the monitoring data store in a table such as a “db_poll_list” table. To add, remove, or update the status of an instance in the monitoring data store, the event processors expose an “updateHost” API that accepts parameters such as a data store identifier, data instance related parameters (e.g., an instance identifier and a DNS address), and an instance status (e.g., “add”, “remove”, or “update”). When an event processor receives this call, the event processor makes the appropriate changes (e.g., adding, removing, or updating an entry) to the db_poll_list table. For example, if a customer submits a request to create a data store with a data store id “id1”, the workflow for creating the data store will, upon provisioning the necessary resources and configuring the data store, mark the state of id1 as “available” in the Admin data store. As a final step in the create database workflow task, the updateHost API can be invoked at one of the event processors, such as by reaching through an internal virtual IP, to add the data store (and its instances) to the monitoring workflow. By making the updating of monitoring status the final (or at least near-final) step in the provisioning workflow, the availability of the creation, deletion, or modification of an RDS data store is decoupled from the availability of the monitoring data store.

Once the host manager sets the status for an active instance to be monitored, the responsible event processor can periodically ping the host manager for the instance as discussed elsewhere herein. If an instance is unavailable, such as may be due to a host machine crashing or rebooting, the event processor will not get a response for the instance and will write information for the potential problem to the Admin data store. A

sweeper will detect the information, and will cause an appropriate recovery workflow to be generated and executed. In one embodiment, a recovery workflow first examines the history of metrics for a data store or data instance, such as information detailing a history of I/O errors for an instance. The workflow then attempts to automatically determine whether the instance is down, such as where there are connection errors, or whether there are no connection problems but an increased number of I/O errors, indicating a potential problem with a particular volume supporting the instance. The tasks of the workflow can attempt to automatically determine and/or isolate the problem, where there are a number of different problems that can occur for a number of different components. Such a determination, as well as the recovery from such problems, is not a trivial matter.

There can be situations, however, where it might not be desirable to automatically recover from a failure. For example, it is possible for an entire data center to fail, where thousands of data stores become unavailable. It can be undesirable to attempt to recover all these data stores at substantially the same time. In one embodiment, the sweeper (or another component of the control plane) can be configured with a maximum number of errors or concurrently executing workflows of a particular type. If a number of workflows exceeds a specified number or threshold, for example, a message or other such notification can be sent or otherwise generated for an operator or DBA, whereby an experienced user can determine the best approach to solving the situation. In one embodiment, the sweeper will run at most a specified number of workflows of the same type at any given time, such as ten workflows of a given type, but will not generate an alarm until a second number, such as twenty-five, or workflows of the same type are requested. A system in accordance with one embodiment provides an operational service dashboard where a DBA or other authorized operator can evaluate the state of the monitoring process(es), and can manually execute recovery actions. Using such an interface, a DBA can select options that enable kicking off workflows, as discussed herein, to perform specific recovery actions. The interface can be used with the control plane to work with multiple disparate database engines and systems, even though the control plane is not in the data path of the data plane. The control plane can monitor error messages and logs, for example, for each of the engines. Such an approach also can allow each data store to be monitored as a whole, concurrently monitoring any replicas of the data store. Different recovery then can be performed based upon the state of the replicas, etc.

It should be recognized that there can be a variety of types of failures that can result in the unavailability or unreliability of a data store or data instance. For example, a host device might fail or reboot, or there might be a problem with the host manager application managing the instance. There also can be a problem with the data store, such as a core dump or segmentation violation (SegV) exception. There also can be problems with the I/O operations or communication paths, or failure of the instance hosting the data store. There also can be various other types of failure, such as failure of a logical volume, a network outage, or an availability zone failure. Different workflows can be used to attempt to determine and recover from the different failure types. In one example, the host manager in one embodiment is the gateway to a respective data instance, and failure of this host manager essentially allows for no control over that instance. To address failures such as a Tomcat process running out of memory, a monitoring component of the control plane can ensure that Tomcat is restarted if necessary. The monitoring system can coordinate restarts to avoid unnecessary error or error detection.

Further, as discussed it is not enough to simply detect and recover from a failure, as other factors must be considered, such as the size or scale of the failure. For instance, the recovery action for the failure of a single cloud instance hosting a data store can be substantially different from a recovery action addressing the failure of an entire availability zone. For larger problems, the multiple failures may need to be correlated and analyzed such that the recovery actions do not compound the existing problems by trying to concurrently recover the various instances individually. In some cases, it might be desirable to perform a staged recovery, where not only are the number of concurrent processes limited, but the ordering of the processes can be controlled such that no data is lost and no recovery actions are taken that later will need to be corrected due to subsequent recovery actions. It also can be desirable in some cases to localize the recovery process as much as possible. It can be beneficial in at least some embodiments to address a failure locally in a safe manner, when possible. For instance, local recovery actions for simple failures such as failure of a host manager or a data process can be preferred to an action performed by an Admin stack of the overall RDS system.

There also can be various reasons for a data instance, data store, or I/O process to fail, each of which might require a different recovery action. For example, a data store bug can cause the data store to fail, or at least generate a significant number of read/write errors. A data store or instance also can fail due to overloads, bad blocks, or other such situations. There also can be user-induced errors, such as an improper query that results in crashing the data store. In other cases, a data store log volume might be filled or corrupted. To address these and other types of failure, the data processes can be constantly monitored by from host manager. As discussed, each host manager can have a status monitor component that checks the status of data store or instance, such as by running a get status command (e.g., for MySQL this can take the form of `/bin/mysql_admin status`). The status monitor component can periodically check the status, and if an instance is unavailable then the instance can be restarted or otherwise addressed. If an instance repeatedly becomes unavailable, or experiences other such errors, the status monitor component can stop attempting to correct the error and cause the information to be written to a monitoring or admin data store in the control plane.

To detect data store errors and I/O crashes, the data store error log and/or kernel log can be monitored in some embodiments. Each host manager can run another module that continually scans for certain error types in these two (or other) error logs, and generates the related metrics. For each error type, a pre-defined threshold can be set, beyond which the errors will be sent to an operator for analysis and possible recovery.

A failure detection mechanism in accordance with one embodiment has a number of constraints applied. For example, it can be configured that the monitoring components scale linearly, such that when the number of data instances exceeds the number of hosts a bucket of event processors are set to poll, for example, additional monitoring components can simply be added as desired. Further, it can be established that all data instances are to be monitored constantly, such as by partitioning the data instances and assigning the monitoring ownership of each partition to one of the event processors. As discussed, if an event processor fails for any reason, the partitions owned and monitored by the failed event processor can be redistributed evenly to other available event processors, such as processors in the same bucket. Also, a list of

database instances can be kept up-to-date by adding tasks to workflows as RDS customers create and delete data stores and/or instances.

Database Recovery Example

As discussed, event processors can check the status of a data instance by making a `getStatus()` (or similar) call to the appropriate host manager. Based on the result of the call, the event processors can determine the appropriate action. As noted earlier, if an event processor cannot connect to the data instance due to an instance failure, for example, then the event processor can move the status of data instance to “incommunicado” and eventually to “dead”. The event processor maintains the state of a data instance in memory. If an event processor fails and a new event processor takes over the monitoring of a “dead” data instance, the new event processor will soon reach the same state independently. One reason for not persisting the state of the data instance is to not depend on the monitoring database for recovery.

Event processors can initiate the recovery of a data instance by changing the status of the data instance in the Admin data store to a state such as “need recovery”, along with a timestamp of the recovery action. An event processor may not be able to mark a data instance for recovery if the database instance status is already marked “need recovery” or “recovering”, for example.

To initiate the recovery action, an event processor detecting a failure marks the failure to the Admin data store. The sweeper component can detect the database marked as “need recovery” and start the recovery process. The sweeper can wait a short amount of time before the recovery starts, in case the cause of the failure is a large scale event requiring human assistance, for example. A threshold exists to limit the number of automatic recoveries that can occur before such a request to call for assistance. Various algorithms or approaches can be used when an event causes a number of instances to be unavailable. For example, if the number of “dead” instances detected is less than a certain threshold (e.g., less than five per period), then a workflow can be kickstarted to start recovering the “dead” data instance(s) individually. This can be accomplished by setting the status of data instance in the Admin data store to a state such as “need recovery” (as long as the state is not already set to a state such as “recovering” or “available”). However, if there are an excessive number of “dead” instances in a given period, this can indicate a correlated or large scale failure. In this case, a notification can be generated for a human operator who can judge best course of action, such as to watch for other events or begin a staged recovery.

A workflow can mark all the recovery-related tasks to be of a higher priority than tasks for creating new data instances or for maintenance tasks. Such an approach allows for the scheduling of recovery-related workflow tasks immediately over other creation or scheduled maintenance-related workflow tasks.

Instance Recovery Example

When an event processor polls a data instance, that instance may be reachable but might not be healthy for reasons such as the host manager having lost all credentials due to failure of the application server, or the instance generating excessive errors due to block storage errors. For both these cases, there is no connection error and the instance is still reachable. If the host manager lost the credentials, the event processor can supply the missing credentials to the host manager as part of the normal operation.

If, however, an instance is generating an excessive number of errors such to failure of the log volumes or data volumes, a specific workflow approach can be used for recovery. For example, a host manager using log analyzers can identify that

data volumes are corrupted. The data files can be recovered using a workflow that obtains the latest snapshots for the underlying data volumes and creates new volumes from these snapshots. The new volumes can be attached to the instance to be recovered. A host manager `dataFileRecover()` or similar API can be called to initiate recovery at the instance. The host manager can be pinged periodically for completion status. Upon receiving the confirmation from the host manager, a data backup procedure can be initiated. The data store then can be restarted and made available to the users, such as by calling the host manager’s `startDatabase` or similar API.

In one embodiment, each host manager implements a `dataFileRecover` or similar API. The host manager can activate and mount each logical volume manager (LVM) on the device, similar to a data store creation setup procedure. The data store can be started in a restricted node, and the latest sequence number of the logs that were backed up can be retrieved and applied to the data store. The workflow component then can be notified by the host manager when the process is complete.

If a host manager instead determines that one of the log files is corrupted, the host manager can restore the log file by downloading the appropriate log files from a storage service or other appropriate location. As noted, the log files can be constantly backed up to persistent storage. In one embodiment, log files are backed up using one common RDS common account with a bin-logs name for a given bucket bucket. The host manager can maintain information such as a customer-id and a database-id in memory, and can retrieve the appropriate log file automatically when needed.

Data Store Partitioning

As is well known in highly-scalable distributed systems, partitioning within a data store only scales to the limits of the physical system in which the data store system resides. Due to this limitation, it can be desirable up front to structure the system in such a way that the system can scale both within a single data storage system, as well as across many data storage systems. Horizontal partitioning of data across distinct data storage systems can contribute to a highly-scalable system which can handle significant demands on the event storage.

A system in accordance with one embodiment utilizes a `customer_id` as the partition key to partition the data tables, including the list of database instances (`db_poll_list`), the related events (`db-events` table), and the security group events table. It can be advantageous to use a customer identifier over a data store identifier, as some events are not restricted to a single data store and may not even concern a particular data store. For instance, a change in a security group does not directly apply to any data store, but may need to be stored as a customer visible event (i.e., retrievable using a `DescribeEvents` API). Further, a single customer’s events may not grow beyond the storage space of a single data store, as in some embodiments event data is only retained for a limited period of time, such as for fourteen days.

There are a number of ways to handle partitioning of data sets across horizontal data store partitions, such as by using bucket partitioning. Bucket partitioning provides an abstraction layer between the data being partitioned and the partitions where the data is being stored. This abstraction layer allows for easier operational management of partitions, such as the addition of new partitions with a migration of data over time, while still allowing for the application to use a hashing mechanism for determining the placement of partitioned data. The implementation of the bucket partition system as described herein comprises components that are specific to

certain embodiments, but the overall concept is applicable to many different use cases as should be apparent.

To implement bucket partitioning, a fixed number of buckets can be determined which are to be available to an application. The number of buckets can remain fixed over the life of the application, such that choosing a large enough number can be important in certain embodiments. The number of buckets can reflect an ability to evenly distribute load across all buckets, which can be individually assigned to a smaller number of physical partitions. If there are too many individual instances assigned to the same bucket, then it can become problematic to efficiently store multiple buckets in a single partition. The fixed number of buckets can act as a middle layer between the data to be partitioned and the partitions themselves. A first step in the layering is figuring out how different pieces of data map to the various buckets. As mentioned above, the partition key for the data can be the customer identifier. An efficient and consistent hashing algorithm can be used to provide a value that can be assigned directly to an individual bucket. Whenever a customer identifier hashes to a value assigned to a bucket, that identifier can live in that bucket for the lifetime of the data.

In this example, buckets are assigned to individual workload partitions. There can always be more buckets than partitions, so a mapping can be used to assign many different buckets to individual partitions. To make the assignment configuration concise, ranges of the bucket numbers can be used to assign the buckets to individual partitions. The following illustrates an example table showing how the partitioning assignment can work:

Partition 1 = {1-25000}
Partition 2 = {25001-50000}

In this example, bucket numbers 1 through 25,000 are assigned to “Partition 1” while bucket numbers 25,001 through 50,000 are assigned to “Partition 2.” Whenever data needs to be added to the system and the hash of the customer identifier maps the workflow instance to bucket **100**, for example, any data related to that customer (including data stores and security groups) can be inserted into tables which physically live in “Partition 1.” Such an approach also can be used to read any information regarding a customer’s database or security groups, where a request for the events for a given customer whose identifier hashes to bucket 100 will be read from “Partition 1.”

The above example deals with a relatively simple case, with the initial assignment of buckets to partitions being unchanged. Sometimes, however, a new partition will need to be added to the system to alleviate the burden on the other partitions. Using this example above, a new partition “Partition 3” can be added to take load off of the other two partitions:

Partition 1 = {1-16666}
Partition 2 = {33333-50000}
Partition 3 = {16667-33333}

As can be seen, 8334 buckets (numbers 16667 through 25000) have been taken from “Partition 1” and re-assigned to “Partition 3.” Also, 8333 additional buckets (numbers 25001 through 33333) have been taken from “Partition 2” and re-assigned to “Partition 3.” This reassignment could have been

based on the buckets which were most busy or most full, but in this example there was a relatively even redistribution of buckets across partitions.

As the bucket assignment changes, the data residing in the physical partition can be affected. In an example above, bucket 100 was used to store the information for a customer whose identifier hashed to 100. In this repartitioning scenario, the data would not be affected since bucket 100 stayed on “Partition 1.” There may have been data in bucket 11000, however, and any data written prior to the repartitioning lives in “Partition 1”, but any data written after the repartitioning will exist in “Partition 3”. To resolve this issue with previous data existing in one partition and current data existing in another partition, the system can allow for more than one partition to be assigned to a bucket. A given bucket can have at least two partitions, a current partition and a previous partition. In the present example, the repartitioning would result in buckets 10001 through 15000 having two partitions assigned, with “Partition 3” as the current partition, and “Partition 1” as the previous partition. As mentioned, any new data for bucket 11000 will be in the current partition, while any data written prior to repartitioning will be in the previous partition. When a query for events or any information maps to bucket 11000, it can be important to check the current partition for that data, as well as to check the previous partition since the record could exist there as well. Such support for multiple partition lookups in a bucket can incur the potential cost of misses for those instances which end up in the previous partition for a given bucket. Since any newly created events are being written to the current partition, however, the cost of a miss will only be incurred for workflow instances running when the repartitioning happens or for closed workflows. Favoring newly created events can improve performance while still allowing the flexibility to do repartitioning efficiently.

As discussed previously, the use of a control plane or service in accordance with various embodiments does not restrict the type of SQL queries that a customer can run, and does not impose any restrictions relating to construction of a schema, such as to be partition ready and not allow queries spanning partitions. Instead, a repository such as a relational database can be provisioned in a computing cloud without restricting the users’ schema or queries. As commonly known, even though there is a theoretical SQL standard, the SQL quirks, syntaxes and their behaviors (e.g., NULL handling) vary across different relational database engines (e.g., MySQL, Oracle, or Postgres). For at least these reasons, users may wish to choose a relational database engine that is familiar for purposes of programming and operations. Such an approach allows customers to use the same set of database tools that the customers have used previously for tasks such as data modeling, development, and debugging, even when the customers migrate their data stores to the cloud (or elsewhere) via the control plane. Using such an approach, customers are not required to rewrite their application or any operational tools, which lowers the barrier of entry significantly for customers to move data to the cloud.

A customer’s data repositories can be moved to the cloud in one embodiment by running the repositories on compute nodes of a cloud computing environment. Block level storage volumes, such as off-instance storage volumes that persist independently from the life of an instance, can be used with these instances for storing the repository binary, logs and volumes, for example. Such an approach can be advantageous, as the virtualization provides flexibility to quickly and

easily scale a compute and storage resources for a repository. Further, such an approach can provide for persistent storage in the cloud.

As known in the art, relational databases can be run in different modes, such as may include: stand-alone (non-replicated), replicated, or replicated and partitioned. A customer typically makes the choice of which mode to run for a repository based on the availability and scalability needs of the repository and the incurred total cost of ownership (TCO). Some applications and services do not require a repository to be highly available and durable, and may instead utilize a stand-alone repository that is able to tolerate outages on the order of minutes. Other applications and servers can require a repository to be always available, and require the repository to never lose data even in the event of a failure. In this case, the applications and services typically require a replicated database offering. Some users, applications, or services require a massively scalable repository that can partition data across multiple repositories, such that scaling can occur beyond the compute and storage capacity of a single database. To address these different use cases, an approach in accordance with one embodiment offers at least two modes, such as stand-alone and high availability, for each database engine. Some embodiments also allow customers build their own partitioning layer on top of either stand-alone or high availability repositories.

As mentioned, the control plane layer can take advantage, or “sit on top,” of various basic software frameworks for performing tasks such as: implementing workflows, establishing secure communication channels between the host managers of the data plane and the components of the control plane, installing software on the instances of the data plane, and performing various database backup and recovery procedures.

For example, a control plane layer can take advantage of a workflow service to manage workflows. As commonly known, a key characteristic of any workflow engine is that the engine enables asynchronous and resumable processing. As discussed above, a workflow can be thought of as a state machine that starts with an initial state and goes through a series of intermediate state transitions by executing different steps of the workflow before reaching the end goal. This end goal can be thought of as the terminal state of a state machine. A workflow service offers the ability to create workflows, and provides hooks to determine the current state of a given workflow and the step(s) to next be executed. The service can store the current state of the state machine, keeping track of the steps which executed successfully and the steps that must be executed to keep the workflow moving. The service does not, in general, actually execute the state transitions for us. The precise tasks of executing the tasks for a workflow will in many embodiments be performed by the “client” components of the workflow.

As discussed above, the various embodiments can be implemented in a wide variety of operating environments, which in some cases can include one or more user computers, computing devices, or processing devices which can be used to operate any of a number of applications. User or client devices can include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless, and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system also can include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These

devices also can include other electronic devices, such as dummy terminals, thin-clients, gaming systems, and other devices capable of communicating via a network.

Various aspects also can be implemented as part of at least one service or Web service, such as may be part of a service-oriented architecture. Services such as Web services can communicate using any appropriate type of messaging, such as by using messages in extensible markup language (XML) format and exchanged using an appropriate protocol such as SOAP (derived from the “Simple Object Access Protocol”). Processes provided or executed by such services can be written in any appropriate language, such as the Web Services Description Language (WSDL). Using a language such as WSDL allows for functionality such as the automated generation of client-side code in various SOAP frameworks.

Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as TCP/IP, OSI, FTP, UPnP, NFS, CIFS, and AppleTalk. The network can be, for example, a local area network, a wide-area network, a virtual private network, the Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network, and any combination thereof.

In embodiments utilizing a Web server, the Web server can run any of a variety of server or mid-tier applications, including HTTP servers, FTP servers, CGI servers, data servers, Java servers, and business application servers. The server(s) also may be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java®, C, C# or C++, or any scripting language, such as Perl, Python, or TCL, as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle®, Microsoft®, Sybase®, and IBM®.

The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (“SAN”) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers, or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (e.g., a mouse, keyboard, controller, touch screen, or keypad), and at least one output device (e.g., a display device, printer, or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices, and solid-state storage devices such as random access memory (“RAM”) or read-only memory (“ROM”), as well as removable media devices, memory cards, flash cards, etc.

Such devices also can include a computer-readable storage media reader, a communications device (e.g., a modem, a network card (wireless or wired), an infrared communication device, etc.), and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed, and/or removable storage devices as well as storage media for temporarily

25

and/or more permanently containing, storing, transmitting, and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services, or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or Web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including portable software, such as applets), or both. Further, connection to other computing devices such as network input/output devices may be employed.

Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules, or other data, including RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

What is claimed is:

1. A computer-implemented method of monitoring components in a data environment, comprising:

under control of one or more computer systems configured with executable instructions,

determining a set of event processors in a control environment for monitoring a plurality of components in the data environment, the plurality of components each having an identifier over a range of identifiers; allocating a portion of the range of identifiers to each of the set of event processors, each event processor being allocated a substantially equivalent portion of the range of identifiers for monitoring;

periodically sending a status message from each of the event processors to be received by the other event processors in the set indicating that the event processor sending the status message is active; and

in response to not receiving a status message from one of the event processors for at least a determined period of time, automatically reallocating the range of identifiers to the active event processors from which status messages were received,

wherein each active event processor receives a different substantially equivalent portion of the range of identifiers based on the number of active event processors, and wherein each event processor is operable to periodically send status messages to each component in

26

the data environment having an identifier in the portion of the range of identifiers allocated to the event processor.

2. A computer-implemented method according to claim 1, wherein:

each event processor is further operable to store information for a monitored component to a job queue when a potential problem is detected, the information being used to determine a recovery workflow to be executed for the monitored component in the data environment.

3. A computer-implemented method according to claim 1, further comprising:

storing information for the event processor from which a status message was not received to a job queue in the control environment; and

using the information to generate a workflow to restart the event processor or add a new event processor to the set of event processors.

4. A computer-implemented method according to claim 3, further comprising:

causing the newly started event processor to send and receive heartbeats to other event processors in the set before activating the newly started event processor.

5. A computer-implemented method according to claim 3, further comprising:

when the newly started event processor is activated, reallocating the portion of the range of identifiers to each of the set of event processors.

6. A computer-implemented method according to claim 1, further comprising:

sorting the identifiers and allocating the sorted identifiers substantially uniformly across the set of event processors.

7. A computer-implemented method according to claim 1, further comprising:

running each event processor at a portion of a processing capacity, wherein each event processor is able to accept a portion of the workload from an unavailable event processor without negatively impacting performance.

8. A system for monitoring components in a data environment, comprising:

at least one processor; and

memory including instructions that, when executed by the at least one processor, cause the system to:

determine a set of event processors in a control environment for monitoring a plurality of components in the data environment, the plurality of components each having an identifier over a range of identifiers;

allocate a portion of the range of identifiers to each of the set of event processors, each event processor being allocated a substantially equivalent portion of the range of identifiers for monitoring;

periodically send a status message from each of the event processors to be received by the other event processors in the set indicating that the event processor sending the status message is active;

in response to not receiving a status message from one of the event processors for at least a determined period of time, automatically reallocate the range of identifiers to the active event processors from which status messages were received;

store information for the event processor from which a status message was not received to a job queue in the control environment; and

use the information to generate a workflow to restart the event processor or add a new event processor to the set of processors,

27

wherein each active event processor receives a different substantially equivalent portion of the range of identifiers based on the number of active event processors.

9. A system according to claim 8, wherein the instructions, when executed by the at least one processor, further cause the system to:

when the newly started event processor is activated, reallocate the portion of the range of identifiers to each of the set of event processors.

10. A system according to claim 8, wherein the instructions, when executed by the at least one processor, further cause the system to:

sort the identifiers and allocate the sorted identifiers substantially uniformly across the set of event processors.

11. A computer-implemented method of monitoring components in a data environment, comprising:

under control of one or more computer systems configured with executable instructions,

determining a set of event processors in a control environment for monitoring a plurality of components in the data environment, the plurality of components each having an identifier over a range of identifiers;

allocating a portion of the range of identifiers to each of the set of event processors, each event processor being allocated a substantially equivalent portion of the range of identifiers for monitoring;

periodically sending a status message from each of the event processors to be received by the other event processors in the set indicating that the event processor sending the status message is active;

in response to not receiving a status message from one of the event processors for at least a determined period of time, automatically reallocating the range of identifiers to the active event processors from which status messages were received;

storing information for the event processor from which a status message was not received to a job queue in the control environment; and

using the information to generate a workflow to restart the event processor or add a new event processor to the set of event processors,

wherein each active event processor receives a different substantially equivalent portion of the range of identifiers based on the number of active event processors.

12. The computer-implemented method of claim 11, further comprising:

causing the newly started event processor to send and receive heartbeats to other event processors in the set before activating the newly started event processor.

13. The computer-implemented method of claim 11, further comprising:

when the newly started event processor is activated, reallocating the portion of the range of identifiers to each of the set of event processors.

14. The computer-implemented method of claim 11, further comprising:

sorting the identifiers and allocating the sorted identifiers substantially uniformly across the set of event processors.

15. The computer-implemented method of claim 11, further comprising:

running each event processor at a portion of a processing capacity, wherein each event processor is able to accept a portion of the workload from an unavailable event processor without negatively impacting performance.

28

16. A computer-implemented method of monitoring components in a data environment, comprising:

under control of one or more computer systems configured with executable instructions,

determining a set of event processors in a control environment for monitoring a plurality of components in the data environment, the plurality of components each having an identifier over a range of identifiers;

allocating a portion of the range of identifiers to each of the set of event processors, each event processor being allocated a substantially equivalent portion of the range of identifiers for monitoring;

sorting the identifiers and allocating the sorted identifiers substantially uniformly across the set of event processors;

periodically sending a status message from each of the event processors to be received by the other event processors in the set indicating that the event processor sending the status message is active; and

in response to not receiving a status message from one of the event processors for at least a determined period of time, automatically reallocating the range of identifiers to the active event processors from which status messages were received,

wherein each active event processor receives a different substantially equivalent portion of the range of identifiers based on the number of active event processors.

17. A computer-implemented method of monitoring components in a data environment, comprising:

under control of one or more computer systems configured with executable instructions,

determining a set of event processors in a control environment for monitoring a plurality of components in the data environment, the plurality of components each having an identifier over a range of identifiers;

allocating a portion of the range of identifiers to each of the set of event processors, each event processor being allocated a substantially equivalent portion of the range of identifiers for monitoring;

running each event processor at a portion of a processing capacity, wherein each event processor is able to accept a portion of the workload from an unavailable event processor without negatively impacting performance;

periodically sending a status message from each of the event processors to be received by the other event processors in the set indicating that the event processor sending the status message is active; and

in response to not receiving a status message from one of the event processors for at least a determined period of time, automatically reallocating the range of identifiers to the active event processors from which status messages were received,

wherein each active event processor receives a different substantially equivalent portion of the range of identifiers based on the number of active event processors.

18. A system for monitoring components in a data environment, comprising:

at least one processor; and

memory including instructions that, when executed by the at least one processor, cause the system to:

determine a set of event processors in a control environment for monitoring a plurality of components in the data environment, the plurality of components each having an identifier over a range of identifiers;

allocate a portion of the range of identifiers to each of the set of event processors, each event processor being

29

allocated a substantially equivalent portion of the
 range of identifiers for monitoring;
 periodically send a status message from each of the
 event processors to be received by the other event
 processors in the set indicating that the event proces- 5
 sor sending the status message is active;
 in response to not receiving a status message from one of
 the event processors for at least a determined period of
 time, automatically reallocate the range of identifiers
 to the active event processors from which status mes- 10
 sages were received; and
 when a newly started event processor is activated, real-
 locate the portion of the range of identifiers to each of
 the set of event processors;
 wherein each active event processor receives a different 15
 substantially equivalent portion of the range of iden-
 tifiers based on the number of active event processors.
19. A system for monitoring components in a data environ-
 ment, comprising:
 at least one processor; and 20
 memory including instructions that, when executed by the
 at least one processor, cause the system to:
 determine a set of event processors in a control environ-
 ment for monitoring a plurality of components in the

30

data environment, the plurality of components each
 having an identifier over a range of identifiers;
 allocate a portion of the range of identifiers to each of the
 set of event processors, each event processor being
 allocated a substantially equivalent portion of the
 range of identifiers for monitoring;
 sort the identifiers and allocate the sorted identifiers
 substantially uniformly across the set of event proces-
 sors;
 periodically send a status message from each of the
 event processors to be received by the other event
 processors in the set indicating that the event proces-
 sor sending the status message is active; and
 in response to not receiving a status message from one of
 the event processors for at least a determined period of
 time, automatically reallocate the range of identifiers
 to the active event processors from which status mes-
 sages were received,
 wherein each active event processor receives a different
 substantially equivalent portion of the range of identifi-
 ers based on the number of active event processors.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 8,060,792 B2
APPLICATION NO. : 12/415987
DATED : November 15, 2011
INVENTOR(S) : Swaminathan Sivasubramanian et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In Column 26, line 61, in claim 8, delete “:” and replace with “;”.

Signed and Sealed this
Twenty-fourth Day of January, 2012

A handwritten signature in black ink that reads "David J. Kappos". The signature is written in a cursive style with a large initial 'D' and 'K'.

David J. Kappos
Director of the United States Patent and Trademark Office